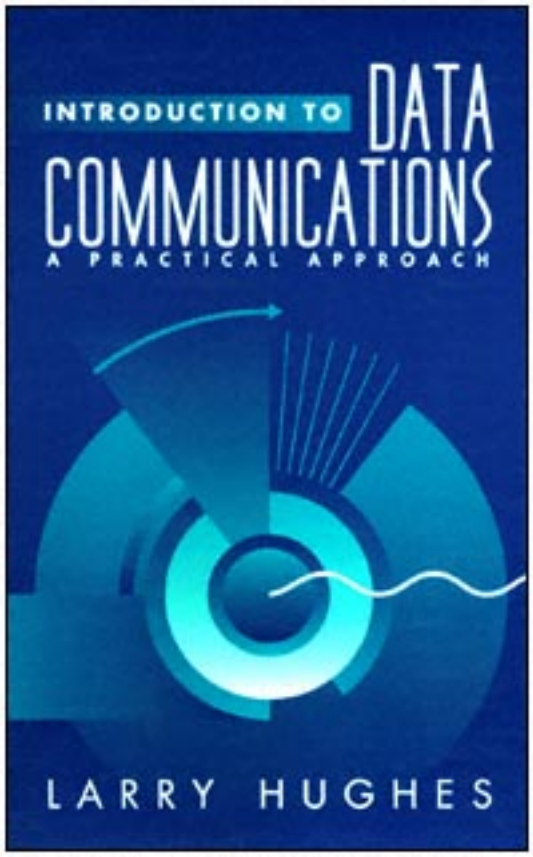

title :	Introduction to Data Communications : A Practical Approach
author :	Hughes, Larry.
publisher :	Jones & Bartlett Publishers, Inc.
isbn10 | asin :	0867207450
print isbn13 :	9780867207453
ebook isbn13 :	9780585364896
language :	English
subject	Data transmission systems.
publication date :	1997
lcc :	TK5105.H83 1997eb
ddc :	004.6
subject :	Data transmission systems.

Introduction to Data Communications

Jones and Bartlett Books in Computer Science

Arthur J. Bernstein and Philip M. Lewis
Concurrency in Programming and Database Systems

Robert L. Causey
Logic, Sets, and Recursion

K. Mani Chandy and Stephen Taylor
An Introduction to Parallel Programming

Harvey G. Cragon
Memory Systems and Pipelined Processors

Nell Dale and Henry Mackay Walker
Abstract Data Types

Nell Dale and John C. McCormick
Ada plus Data Structures

Nell Dale
A Laboratory Course in Turbo Pascal, Second Edition

Nell Dale, Mark Headington, and Chip Weems
Programming and Problem Solving with C++

Michael J. Flynn
Computer Architecture; Pipeline and Parallel System Design

John Gregory and Don Redmond
Introduction to Numerical Analysis

James Hoin
Discrete Structures, Logic and Computability

Larry Hughes
An Introduction to Data Communications

E. Stewart Lee
Algorithms and Data Structures in Computer Engineering

Peter Linz
An Introduction to Formal Languages and Automata, Second Edition

Richard Neapolitan and Kumarss Naimipour
Foundations of Algorithms

Christopher H. Nevison, Daniel C. Hyde, G. Michael Schneider, and Paul T. Tymann, Editors
Laboratories for Parallel Computing

Greg W. Scragg
Problem Solving with Computers

Charles Van Loan
An Introduction to Computational Science and Mathematics

Henry M. Walker
The Limits of Computing

Nell Dale and Susan C. Lilly
Pascal Plus Data Structures, Algorithms and Advanced Programming, Fourth Edition

Nell Dale and Chip Weems
Introduction to Turbo Pascal and Software Design, Fourth Edition

Nell Dale and Chip Weems
Intro to Pascal and Structured Design, Fourth Edition

Mark R. Headington and David D. Riley
Data Abstraction and Structures Using C++

William Ford and William Topp
Assembly Language and Systems Programming for the M68000 Family, Second Edition

Nell Dale, Susan C. Lilly and John McCormick
Ada Plus Data Structures: An Object-Based Approach

Nell Dale, Chip Weems, and John McCormick
Programming and Problem Solving with Ada

James Roberge
Data Structures in C++: A Laboratory Course

James Roberge
Data Structures in Pascal: A Laboratory Course

James Roberge
Introduction to Programming in C++: A Laboratory Course

Philippe J. Gabrini and Barry L. Kurtz
Data Structures and Algorithms with Modula-Z

Yuksel Uckan
Application Programming and File Processing in COBOL: Concepts, Techniques, and Applications

James L. Noyes
Artificial Intelligence with Common LISP: Fundamentals of Symbolic and Numeric Processing

James L. Hein
Theory of Computation: An Introduction

Introduction to Data Communications

A Practical Approach

Larry Hughes

Department of Mathematics and Computing Science
Saint Mary's University
Halifax, Nova Scotia

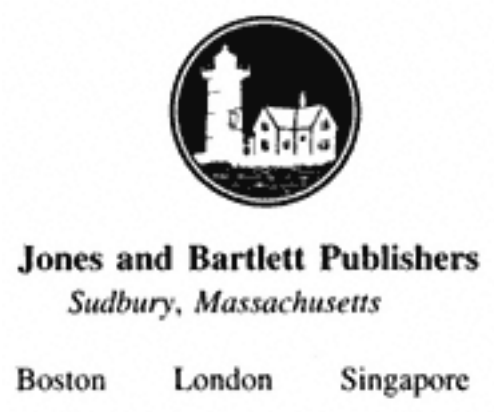

Jones and Bartlett Publishers
Sudbury, Massachusetts

Boston London Singapore

Disclaimer:
This netLibrary eBook does not include the ancillary media that was packaged with the original printed version of the book.

Editorial, Sales, and Customer Service Offices

Jones and Bartlett Publishers
40 Tall Pine Drive
Sudbury, MA 01776
508-443-5000
1-800-832-0034
info@jbpub.com
http://www.jbpub.com

Jones and Bartlett Publishers International
Barb House, Barb Mews
London W6 7PA
UK

Copyright © 1997 by Jones and Bartlett Publishers, Inc.

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilized in any form, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the copyright owner.

Library of Congress Cataloging-in-Publication Data

Hughes, Larry, 1954
Introduction to data communications: a practical approach / Larry Hughes.
p. cm.
Includes bibliographical references and index.
ISBN 0-86720-745-0
1. Data transmission systems. I. Title.
TK5105-.H83 1996 96-13531
004.6dc20 CIP

Acquisitions Editor: Dave Geggis
Production Administrator: Anne S. Noonan
Manufacturing Manager: Dana L. Cerrito
Editorial Production Service: Superscript Editorial Production Services
Typesetting: ICPC
Cover Design: Hannus Design Associates
Printing and Binding: Courier Companies, Inc.
Cover Printing: Coral Graphic Services, Inc.

Printed in the United States of America
00 99 98 97 96 10 9 8 7 6 5 4 3 2 1

CONTENTS

Preface xi

Part 1
Background 1

1
Basic Concepts and Terminology 3

1.1 Introduction 3

1.2 Entities 3

1.3 Protocols 6

1.4 Data Terminal Equipment 16

1.5 Summary 19

1.6 Exercises 19

2
Commkit 21

2.1 Introduction 21

2.2 The 8086 Family and Its Hardware 21

2.3 The Commkit Software 31

2.4 Interprocess Communication (IPC) Using Commkit 51

2.5 Compiling Commkit 54

2.6 Summary 56

2.7 Exercises 56

Part 2
Point-to-Point Communications 59

3
Asynchronous Communications 60

3.1 Introduction 60

3.2 General Concepts 60

3.3 Asynchronous Communication Hardware 69

3.4 The PC UART 76

3.5 Point-to-Point Communication Using Commkit 88

3.6 Example 93

3.7 Summary 99

3.8 Exercises 100

4
Physical Connections 103

4.1 Introduction 103

4.2 Interconnecting DTEs 104

4.3 Standards 110

4.4 RS-232-C (V.24) 111

4.5 Other Interconnection Standards 113

4.6 Summary 117

4.7 Exercises 118

5
Testing Communication Systems 119

5.1 Introduction 119

5.2 Built-In Tools 119

5.3 Channel-Monitoring Tools 122

5.4 The Commkit Line Analyzer 123

5.5 Using the Commkit Line Analyzer 128

5.6 Summary 129

5.7 Exercises 129

Part 3
Longer-Distance Communications 133

6
Synchronous Communications 136

6.1 Introduction 136

6.2 Reducing Transmission Overheads 137

6.3 Improving Error Detection 142

6.4 Error Recovery 154

6.5 Synchronous Protocols 178

6.6 Synchronous Communications Hardware 187

6.7 Commkit Synchronous Software 189

6.8 Summary 196

6.9 Exercises 197

7
File Transfer 200

7.1 Introduction 200

7.2 File Systems 201

7.3 File Transfer Protocols 203

7.4 Commkit File Transfer Example 205

7.5 Summary 209

7.6 Exercises 209

8
The Telephone System 211

8.1 Introduction 211

8.2 The Telephone System 211

8.3 Connecting to the Telephone System 214

8.4 RS-232-C (ITU-T V.24) 221

8.5 High-Speed (Synchronous) Modems 230

8.6 PC Modem Control 239

8.7 Modem Testing 242

8.8 Modem Variations 244

8.9 Commkit External Modem Software 245

8.10 Commkit Internal Modem Software 250

8.11 Summary 253

8.12 Exercises 254

9 Multiplexing 256

9.1 Introduction 256

9.2 Background 256

9.3 Multiplexing Techniques 259

9.4 Port Selectors 262

9.5 Multiplexing with Commkit 265

9.6 Summary 270

9.7 Exercises 271

10
Wide Area Networks 273

10.1 Introduction 273

10.2 Alternatives to the Public Telephone Network 274

10.3 Wide Area Networks 279

10.4 Wide Area Network Protocols 296

10.5 Examples of Wide Area Networks 306

10.6 The Commkit Wide Area Network 306

10.7 Using the Commkit Wide Area Network 319

10.8 Summary 325

10.9 Exercises 325

Part 4
Local Area Networks 329

11
Bus Local Area Networks 331

11.1 Introduction 331

11.2 Examples 335

11.3 The Commkit Bus Network 342

11.4 Using the Commkit Bus Network 352

11.5 Summary 354

11.6 Exercises 354

12
Ring Local Area Networks 358

12.1 Introduction 358

12.2 The Token Bus 358

12.3 Ring Networks 363

12.4 The Commkit Token Ring 376

12.5 Using the Commkit Token Ring 387

12.6 Summary 388

12.7 Exercises 388

Part 5
Internetwork Communications 393

13
Layered Architectures 395

13.1 Introduction 395

13.2 Background 395

13.3 Terminology 399

13.4 Examples of Layered Architectures 401

13.5 Commkit Sockettes 412

13.6 Summary 422

13.7 Exercises 423

14
Interconnecting Networks 426

14.1 Introduction 426

14.2 Background 427

14.3 Repeaters 430

14.4 Bridges 430

14.5 Gateways 437

14.6 The Commkit Transparent Bridge 440

14.7 Summary 450

14.8 Exercises 451

Appendix A
An Introduction to C 455

A.1 Comments 455

A.2 Base Types 455

A.3 Unnamed Constants 455

A.4 Identifiers and Variables 456

A.5 Expressions 458

A.6 The Assignment Statement 459

A.7 Selection 460

A.8 Iteration 462

A.9 Aggregate Data Types 465

A.10 Functions 470

A.11 Complier Directives 473

A.12 Software Management 475

Appendix B
Running and Testing Commkit 477

B.1 General Hints 477

B.2 A Final Thought 478

Appendix C
State Diagrams 479

C.1 Concepts 479

C.2 Implementation 481

Appendix D
Character Codes 483

D.1 Five-Bit Baudot 483

D.2 Seven-Bit ASCII 484

D.3 Eight-Bit EBCDIC 486

Appendix E
Single-Port Operations 488

E.1 Basic Concepts 488

E.2 The Parallel Port 489

E.3 The Line Analyzer 489

E.4 The Bus Network 489

E.5 The Ring Network 490

E.6 The Commkit Bridge 490

E.7 A Final Word 490

Suggested Readings 491

Index 493

PREFACE

A fundamental concept in computing science, from the basic electronic circuit through programming languages to user applicions, is input-processing-output. This concept is often overlooked and taken for granted in communications. Communications have always been an integral part of the development of computers and computing science: from von Neumann's original architecture through the early timesharing systems of the 1960s to the growing array of computer networks of today.

As with most aspects of computing science, communications can be discussed in terms of different levels, or layers. At the highest level is the user who wants to transfer information from one location to another (whether a file from a computer on one side of the world to the user's local computer, or simply a character from a terminal to a central computer), while at the lowest level is information, in the form of bits, moving between computers. The term data communications is usually applied to the "lower layers" of a communicationthat is, those responsible for the transfer of the information (i.e., the data) to the intended destination(s). Since most users are well removed from the intricacies of the system supporting the communication, the lower layers may only become an issue should the system fail. For example, to the user of a remote database communications become important if the communication system fails. However, to the person maintaining the communication system, the information is irrelevant; what is important is its reliable transfer.

Although traditionally an engineering discipline, data communications is rapidly becoming of interest to professional programmers as well as to many university and college computing science programs for a number of reasons, including the following:

The growing demand for computers to support different types of communication.

The development and reliance upon communication software.

The expanding use of computer networks.

With the appropriate facilities, data communications can be taught in a practical, hands-on manner in which, ideally, students will have access to equipment that supports a number of networks and permits experimentation with as many aspects of data communications as possible. However, there are at least two reasons why many data communications courses are *not* taught in this way:

1. Much of the equipment needed to teach data communications is expensive, even when used with low-cost personal computers such as the ubiquitous PC. The costs associated with an Ethernet card for a single PC, for example, can range anywhere from $75 to $150, and additional expenses may be incurred since many local area networks require further specialized hardware and software to manage the network. If a number of different networks are taught, the costs escalate rapidly.

2. Even if networking facilities are available, many network managers are reluctant (often to the point of hostility) to allow unrestricted network access, especially to undergraduates.

Fortunately, it *is* possible to teach data communications without purchasing additional equipment (or incurring the wrath of the network manager) and at the same time allow students to gain a practical understanding of the subject. For example, this book, with its software, covers the major areas of data communications in a practical manner using the PCthe only additional expense being a second serial port (if the machine is not already equipped with one). The software, known as Commkit, has been designed to permit anyone with access to at least two PCs to gain experience using the concepts associated with point-to-point communications, network analysis tools, modems, store-and-forward wide area networks, bus and ring local area networks, gateways, and a version of UNIX sockets.

Commkit has been designed to be as unobstrusive as possible to allow the reader to concentrate on communication issues rather than on the operating system or the hardware. Written entirely in Borland C, the software supports message passing between low-level interrupt handlers, background processes, and a foreground process.

This book begins with an examination of the methods by which communications can take place between two computers separated by distances of less than 50 metersthat is, limited-distance point-to-point communications. Next, techniques for handling longer-distance communications and larger-volume data transfers are presented: first, the telephone system and then wide area networks. Having outlined the principles of networks and synchronous protocols, the book then considers local area network communications. Finally, mechanisms to support communications between computers on separate network are introduced.

The topics covered here are all examined using the Commkit software to highlight specific areas of data communications. All necessary executables, object modules, and source files are supplied on the Commkit diskette for each topic. The executables can be used directly or recreated by modifying the source modules.

This book is divided into five parts. Part 1 introduces the reader to data communications and the Commkit software. In Chapter 1, the basic concepts and terminology associated with communications and data communications in particular are covered. The material presented is used throughout the remainder of the text. Chapter 1 concludes, as do all subsequent ones, with a series of exercises covering the material just presented. Many of the exercises use the executables supplied on the Commkit diskette; others require the modification and recompilation of the source software.

The Commkit software, its operation, and the interaction between Commkit and the PC's hardware are described in Chapter 2. A detailed discussion of Commkit, its design, and internal structure is presented by examining its use in interprocess communication.

The communication that most readers of this book will encounter involves the transmission of information between pairs of objects (such as a PC connected to a central computer), which is known as point-to-point communication. This is the subject of Part 2. Once the terminology and concepts are understood, nearly all other types of communication can be grasped.

Point-to-point communications are introduced in Chapter 3 using the asynchronous communication hardware available on all PCs. The chapter presents a detailed examination of how asynchronous communications are achieved, using Commkit's point-to-point telephone utility that supports communications between pairs of PCs as a detailed case study.

Some of the standards and methods of physically supporting a communication are introduced in Chapter 4. Standards examined include RS-232 and twisted-pair and coaxial cable are among the media considered.

One of the biggest headaches in the detection and correction of errors in a communication system is determining what is actually being sent across the channel between the communicating devices. This is true in the commercial world as well as in a course on data communications. To assist in error detection and correction, as well as to monitor what is taking place on the communication channel, commercial tools known as line analyzers are available. Since commercial line analyzers can be extremely expensive, Commkit is supplied with a software equivalent. Both Commkit's version and commercial line analyzers are discussed in Chapter 5.

Part 3 moves on from communications consisting solely of single-byte transmission between pairs of computers to the transfer of large volumes of information composed of hundreds or thousands of bytes across a computer network. A number of topics are covered, including sending large volumes of information with minimum overhead; file transfer; using the telephone system to support data communications; maximizing channel utilization; and utilizing networks to handle data communications.

Chapter 6 examines techniques for sending information that minimize the overhead associated with asynchronous communications by reducing the amount of control information sent, and simultaneously improving the reliability of the

data transfer. Commkit is supplied with one such data transfer protocol, which is examined in detail, showing how it supports large-volume transfer while also handling the detection and correction of errors.

The transfer of large volumes of information is typified by copying files between machines. In Chapter 7, file access using Commkit is described and a reliable file transfer protocol is designed. The resulting protocol is implemented atop the reliable large-volume transfer protocol examined in Chapter 6.

How the telophone system handles communications over longer distances is considered in Chapter 8. Methods of representing information in the telephone system are discussed in detail. RS-232-C, originally introduced in Chapter 4, is covered in depth to explain how computers can access the telephone network. High-speed modems that use data compression and synchronous communications are presented as alternatives to existing low-speed modems. Two kinds of modems are examined using Commkit's modem software: the first the external modem and the second the internal modem. Both illustrate how access to the telephone network can be achieved.

Chapter 9 is the first of two chapters that consider methods of reducing the overheads associated with using a telephone for communications. It examines how a single channel can be shared or multiplexed among a number of communicating devices. An example of multiplexing is presented using Commkit's statistical multiplexing software.

The first four chapters of Part 3 introduce wide area networks (WAN), which are then examined in Chapter 10. Several WANs are discussed, and the concepts associated with them are applied using the Commkit store-and-forward network.

Part 4 moves away from the "traditional" areas of data communications into local area networks (LANs), examining two of the best known LAN topologies: the bus and the ring. In Chapter 11, bus local area networks are discussed in light of existing commercial LANs such as the Ethernet. The bus network emulator supplied with the Commkit software is examined and illustrates many of the issues associated with bus LANs.

Chapter 12 builds upon the material covered in the previous chapter to explain ring networks and how they operate. It initially considers the token bus and then ring local area networks. The concepts associated with ring LANs are presented through a detailed examination of the Commkit token ring network.

Part 5 considers internetworking, a methodology that allows machines on different, potentially dissimilar, networks to communicate. Two major internetworking issues are examined in detail: the layering of systems and the interconnection of networks. Layering is described in Chapter 13 using Commkit's version of UNIX sockets, known as sockettes. Sockettes are part of a layered architecture that presents a network-independent addressing structure and function atop any of the Commkit networks discussed in Parts 3 and 4. Chapter 14 examines some of the problems with interconnecting networks. It includes examples of existing internetworking standards, some of which are illustrated using the Commkit bridge software.

Five appendices are also included. Appendix A is a brief introduction to C for readers unfamiliar with the language. A detailed description of how to run and test Commkit is given in Appendix B. An explanation of state diagrams and how they can be implemented is presented in Appendix C. Appendix D contains a listing of several well-known character codes, including ASCII and EBCDIC. Appendix E suggests modifications to the software that permit the support of single-port operations.

Exercises

Each exercise in the text is ranked as follows:

(A) Book exercises that do not require the use of the Commkit software.

(B) Exercises that require the use of the Commkit software and one or more PCs.

(C) Exercises that require the modification and recompilation of one or more Commkit modules. At least one PC is required to test the software.

Notation

Unless otherwise indicated, all notation in the book conforms to the following format:

all software is written in C; source code examples found in the book are displayed in `typewriter font`, as are MS-DOS command.

Control characters such as *ETX* and *STX* are shown in italics.

Keyboard input to Commkit is highlighted using Small Capitals.

Acknowledgments

In writing a textbook, an author always benefits from the experience of others; *Introduction to Data Communications* is no exception. Accordingly, I must thank the following for their helpful suggestions on both the book and the Commkit software:

Lee Joudrey of the Department of National Defence, Halifax.

Liz Leboffe of St. John Fisher College, Rochester.

Doug MacGillivary of Bell Canada, Ottawa.

David Powell of STC Submarine Systems, Ltd., London.

The many students who have used Commkit while taking my data communication courses.

Finally, to Sandy Cook, the one person who always knew that the book and the software would be written. Her constant support and encouragement made this book possible.

PART 1
BACKGROUND

After the American Revolution, the one major seaport on the eastern coast of North America to which the British Navy had access was Halifax, in Nova Scotia. Since the British did not want to lose the rest of their North American colonies, it was imperative that Halifax be defended. Accordingly, a form of "early warning system" was installed, consisting of a series of forts stretching from the Atlantic Ocean to Halifax, that could relay information about a ship sighted off the coast to the Citadel (the fortress in Halifax) in about 15 minutes.

Flags were used to communicate between the forts: Whenever a ship was spotted off the coast, this information was encoded as a unique sequence of flags and then passed from one fort to the next until it reached the Citadel. If the officers there were to react to the encoded information, it was important that it be sent and received correctly by each fort up the coast. For example, the results could have been disastrous if a fort identified a ship as a neutral man-of-war only to have the information corrupted down the line to indicate

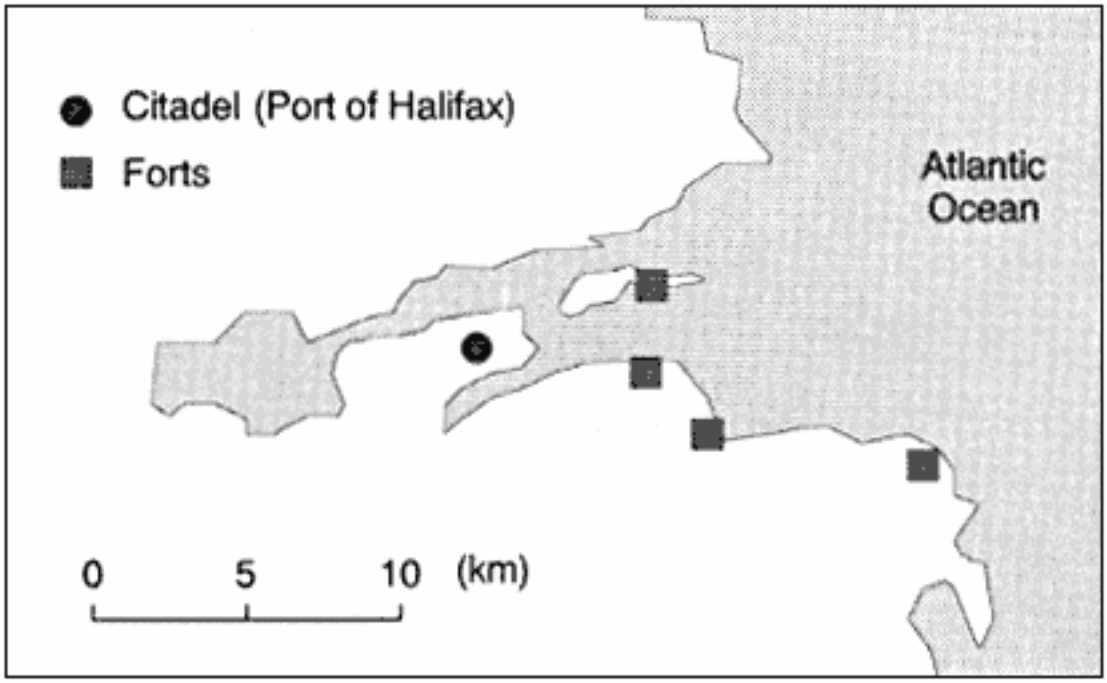

a French man-of-war. Similarly, if an admiral's ship was identified but fog caused the information to be read as "the incoming ship is a scow," the admiral might not have been too pleased with the welcome he received.

This seemingly simple example illustrates all of the major issues surrounding data communications:

Information must be *encoded* before it can be *transmitted*. The sailors watching for incoming ships had to decide among themselves the type of ship and then translate that information into a set of flags.

Rules must be in place to ensure that any information sent is received correctly. At a minimum, both the transmitter and receiver must agree upon the same method of translating the encoded information. In eighteenth century Halifax, all of the forts were required to use the same flags and interpret them in the same way.

When an error is *detected*. there must be an agreed-upon set of procedures to follow that allow the errors to be *corrected*. If errors are permitted to propagate through the system (or, for example, up the chain of command from the officers in the forts to those in the Citadel), the information is of no value and cannot be used or acted upon safely.

Mechanisms must be available to ensure that both the transmitter and receiver remain *synchronized* (or *in step*). Should the communicating parties lose synchronization, information may be lost, with disastrous results (for example, an enemy warship might enter the harbor undetected).

All technical subjects, regardless of the discipline, require their practitioners to understand and be familiar with

The terminology associated with the subject that permits the exchange of ideas and information. Data communications, as a technical subject, has its own terminology, most of which can be explained in terms of everyday activities. Chapter 1 introduces much of the terminology used in data communications both through a series of examples and by considering how humans communicate.

The tools available for working with (or learning about) the subject. Since data communications incorporates hardware as well as software, it is important to have an understanding of both. In Chapter 2, the Commkit software is examined in terms of some of the PC's hardware and the methods by which Commkit supports communications.

1
Basic Concepts and Terminology

1.1
Introduction

Communications, whether between humans, animals, or computers, involves the transfer of information. Considering that people developed computers, it should not be surprising that many of the problems that designers of data communication systems face are, in many respects, similar to those encountered when people communicate on a day-to-day basis. For example:

When two people want to talk, some agreement must be made to ensure that the person who begins talking eventually stops so that the other person has a chance to respond.

When a person misses part of what is said in a conversation, mechanisms should exist that allow the lost information to be repeated.

When a person finishes speaking, it is often considered polite (and sometimes necessary) to ask whether what has been said has been heard and comprehended.

In this chapter, some of the general terms and concepts associated with communications are introduced. In all cases, everyday examples are used to reinforce understanding.

1.2
Entities

Communications, unless otherwise stated, are assumed to be between pairs of objects, often called entities. The transfer of information occurs across a channel (sometimes referred to as a line), of which there are two types:

A simplex channel is one in which communication can occur in one direction only.

A duplex channel is one in which a communication can occur in both directions simultaneously.

The channel can be any medium, including air, copper wire, or optical fibers. At any moment, an entity can be

A transmitterthat is, it sends information.

A receiverthat is, it receives information.

Both a transmitter and a receiverthat is, it can transmit and receive information simultaneously.

Two broad categories describe the type of communication that can take place between entities.

The first type of communication permits either entity to transmit, but not simultaneously, and is known as half-duplex. A typical half-duplex communication involves one entity transmitting its information while the other receives. The roles are then reversed, and the entity that was originally receiving now transmits (while the original transmitter receives). As shown in Figure 1.1, the reversing of roles is known as line turnaround. (In this and subsequent diagrams, *T* denotes the transmitting entity and *R* denotes the receiving entity.)

An example of a half-duplex communication is a "telephone" consisting of two tin cans connected by a piece of string. At any point, only one person can speak (i.e., transmit) while the other listens (i.e., receives). Line turnaround requires the transmitter to inform the receiver that there is no more information to be transmitted and that the receiver can now transmit.

A half-duplex communication can be implemented using a pair of simplex channels (one for each direction) or one duplex channel. In either case, both entities must agree to a mechanism whereby line turnaround can be signaled and their roles reversed.

The second category of communication permits either entity to transmit and receive simultaneously and is known as full-duplex. Full-duplex communications are generally used when both entities have large volumes of data to send or when the transmitter wants to determine if the information sent has been received correctly (without having to perform a line turnaround). A representation of a full-duplex communication is given in Figure 1.2.

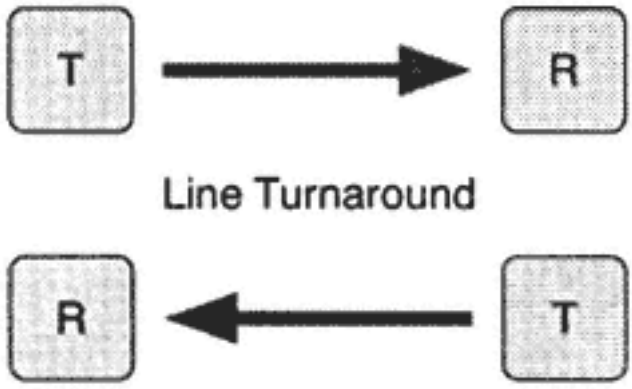

Figure 1.1
Half-duplex communication
and line turnaround.

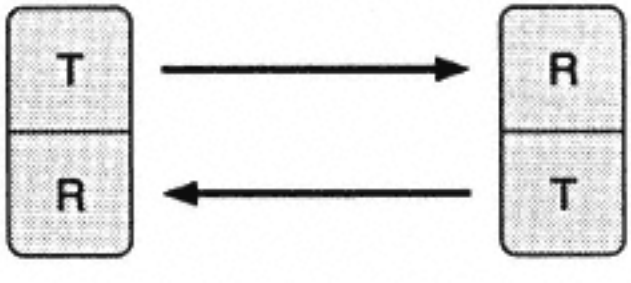

Figure 1.2
A full-duplex communication.

Examples of full-duplex communications include sending mail (a person might send a letter to a friend and receive one from the same friend at the same time) and talking to someone (both people might attempt to talk at the same time).

Full-duplex communications can be realized in a number of ways. If a duplex channel is used, each entity is allocated a distinct part of the channel (for example, in Figure 1.3 one entity may transmit on frequency *X* and receive on frequency *Y*. while the other would transmit on frequency *Y* and receive on frequency *X*).

It is possible to support a full-duplex communication using a pair of simplex channels. In this situation, as illustrated in Figure 1.4, each entity transmits on one channel and receives on the other.

In a communication involving humans, it is obvious who the entities are: the people involved in the communication. However, in a communication involving computers, it is not so easy to determine the entities, since one may be the application process (the software requiring the communication), the support software (the software supporting the communication requirements of the application processes, perhaps needing the communication facilities offered by the underlying hardware), or the processor (computer) itself.

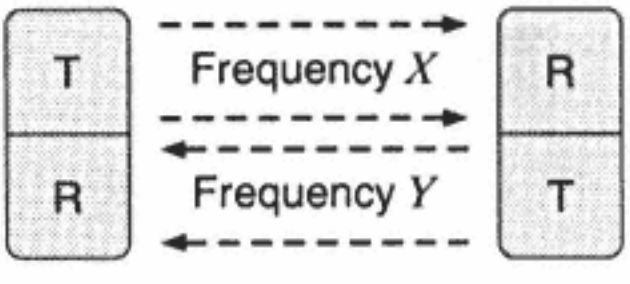

Figure 1.3
Implementing full-duplex communication with
different frequencies on a single channel.

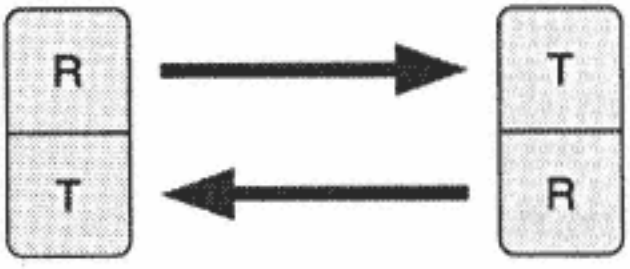

Figure 1.4
Implementing full-duplex communication
using a pair of simplex channels.

For example, a communication between two processes may involve one or more processors. In the single-processor case, only the processor's communication software is necessary, but if the processes reside on separate, interconnected processors, both communication software and hardware are needed.

The structuring of software so that one level uses the services of an adjacent, lower level is known as layering. Layering will be discussed in subsequent chapters.

1.3 Protocols

At the start of this chapter, examples of communications between humans were presented. These examples illustrate the need for a set of rules to ensure that the communication will proceed in spite of such problems as people attempting to speak simultaneously and information being misunderstood or missed entirely.

Other communicating entities, such as computers or processes, also require rules, or protocols, to ensure that the communication can proceed. Protocols are intended both to control the communication between the stations and to define certain characteristics about it. Regardless of what the entity is, all entities participating in the communication must agree on the protocol or the communication may fail. A full-duplex protocol will probably fail, for example, if the communicating entities are connected by a single simplex channel.

Broadly speaking, any communication protocol can be discussed in terms of the following:

Coding of information: how the information is represented between the various entities.

Control: how the communication is controlled by the entities involved in the communication.

Error checking and recovery: how the entities ensure that the information is sent and received correctly.

Channel utilization: how efficiently the channel is used by the communicating entities.

Synchronization and timing: how the entities remain in step during the progress of a communication.

Transparency: how the mechanisms supporting the communication are hidden from the entities.

1.3.1 Coding of Information

For a communication to take place, both the transmitting and receiving entities must agree upon a common language, written or verbal. Thus, for a

person in England to carry on a conversation with a person in France, either the person in England must know how to speak French, or the person in France must know how to speak English, or both must agree on a third language.

Although there is no universal standard for the coding of information in computers, all manufacturers agree upon the value and structure of the bit: a single cell having the value of zero or one. Bits can be grouped into 4-bit quantities known as nibbles and 8-bit quantities known as bytes. Figure 1.5 shows the organization of these three structures.

The description of the byte data structure depends entirely upon the manufacturer, although some standards do exist (see below). For example, some manufacturers label bytes from 1 to 8, while others label bytes from 0 to 7. Similarly, the labeling may run in ascending order from right to left (i.e., from 0 to 7, the rightmost bit being 0, and the leftmost, 7), or in descending order from right to left (the rightmost bit being 7, and the leftmost, 0).

The value and interpretation of the byte depends, once again, on the manufacturer. If bytes are considered unsigned, their possible values are 0 through `255`; however, if they are signed, the range of values is `-128` through `+127`.

Bytes are also used to represent characters: The two most common character codes used by computer manufacturers are

EBCDIC (Extended Binary Coded Decimal Interchange Code), a code made popular by IBM in its large mainframe computers in the 1960s.

ASCII (American Standard Code for Information Interchange), an alternate standard adopted by almost all non-IBM computer manufacturers and now so widespread that even IBM uses it for its PCs.

Appendix C contains character code tables for the 5-bit Baudot code, 7-bit ASCII, and 8-bit EBCDIC.

Not all bytes are intended to represent textual, printable characters; some are used to convey control information such as line feed, carriage return, and form feed. In the table of 7-bit ASCII bytes, for example, the first two columns contain control characters, while the remaining six (with the excep-

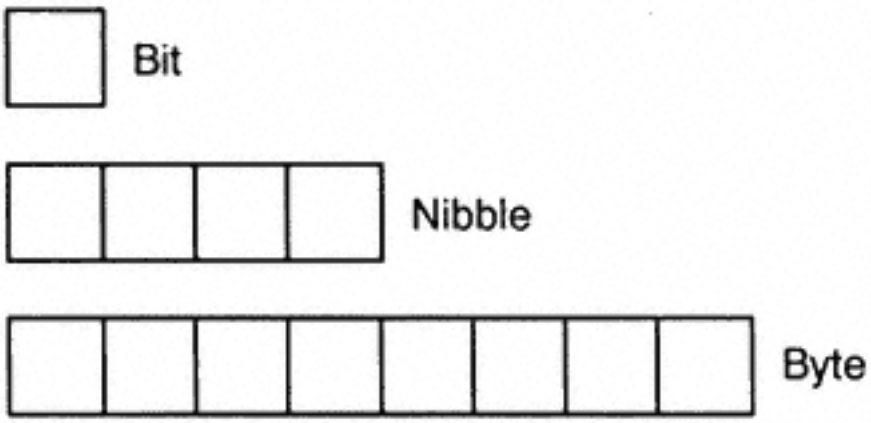

Figure 1.5
The bit, nibble, and byte structures.

tion of *DEL* at the bottom right of the table) represent graphic (or printable) characters.

Although most communications usually occur at the bit or byte level, applications may deal with other data structures that comprise groupings of bytes. Depending upon the manufacturer, a word may be a 16-bit (2-byte) or a 32-bit (4-byte) data structure. Floating-point numbers, consisting of a mantissa (the fractional part) and an exponent, may be stored and interpreted in a variety of ways, depending upon the precision offered by the manufacturer and the number of bits available to represent them (floating-point numbers can be 32-, 64-, or even 128-bit quantities). Fixed-point, or Binary Coded Decimal (BCD), numbers can be stored in different ways, including:

Zoned decimal format has each digit in the BCD number occupying a single byte. The leftmost nibble (zone nibble) contains `1111`, while the rightmost contains the digit (one of `0000` through `1001`). The sign (one of `1100`positive; `1101`negative; or `1111`unsigned, assumed positive) is stored in the leftmost nibble of the low-order digit. Figure 1.6 illustrates how `-218` would be stored.

Packed decimal format removes the zone nibble from each byte and stores each digit in adjacent nibbles. The sign nibble is the same as for zoned decimal, but is stored in the rightmost nibble of the low-order byte. The internal representation of `-218` stored in packed decimal format is shown in Figure 1.7.

How information is represented on different computers is an important issue in data communicationsunless the entities involved in the communication can agree upon how the information is coded, there can be no meaningful exchange. For example, if a file of ASCII information is transmitted to an EBCDIC computer, unless some form of translation is made available (i.e., the ASCII bytes are translated into EBCDIC, or the EBCDIC computer can interpret ASCII), the information received by the EBCDIC computer may be meaningless.

To make matters even more confusing, although two manufacturers may agree upon the character code, the internal organization, or byte ordering,

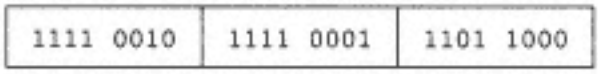

Figure 1.6
Storing `218` as a zoned decimal.

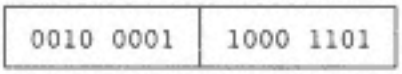

Figure 1.7
Storing `218` as a packed decimal.

of the computer's memory may differ from machine to machine. Figure 1.8 shows how the string ABCDEF is stored in a processor in the Intel 8086 family, whereas Figure 1.9 shows the same string ABCDEF stored in a Motorola 68000 processor with a different byte ordering (but still ASCII).

If data is transmitted between an Intel and a Motorola processor as a series of 8-bit ASCII bytes, the hardware ensures that each byte is read from and written to memory in its proper order. However, if the information is sent as a series of 16-bit words, the receiving processor will receive the bytes in the reverse order. Similar problems can occur when attempting to transmit 16-bit integers as pairs of bytes. In situations such as these, a common byte ordering must be agreed upon (this is discussed in more detail in Chapter 13).

1.3.2
Control

When two people talk, they use many techniques and cues to signal whether the information has been understood, if the information should be repeated, and when each can speak. These signals are not usually part of the information being conveyed, but they are necessary to allow the communication to continue. For example, to gain a student's attention, an instructor could prefix a question with the student's name; however, the name has no bearing on the actual information being conveyed. Other types of such information exist as well: in order to determine if the receiver has actually received (and understood) the information, the speaker may end what has been said with a question, expecting the receiver to respond to it. The additional information used to control the communication is known as control information. (To distinguish between the information being transmitted and the control infor-

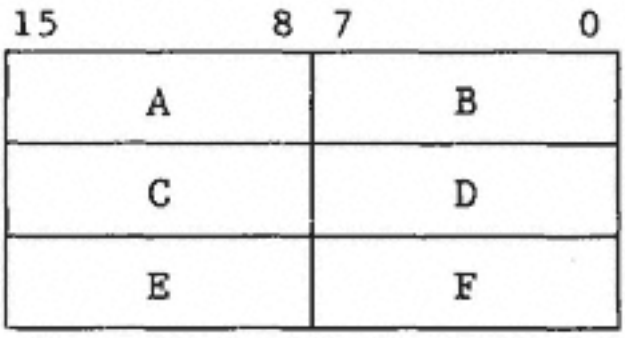

Figure 1.8
Byte ordering in an Intel
8086 family processor.

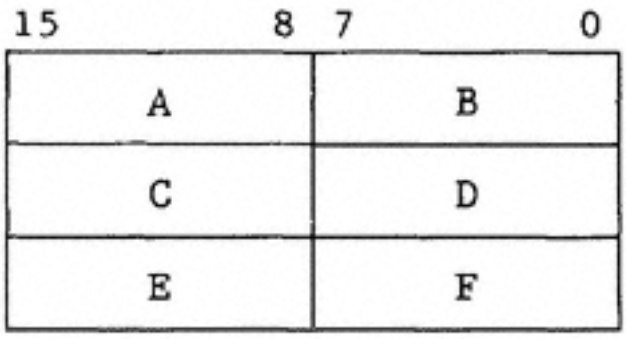

Figure 1.9
Byte ordering in a Motorola
68000 processor.

mation, the non-control information is often referred to as a message.) The extra information required to ensure that the communication can continue is known as an overhead. It is intended to ensure that the communication can proceed and that the message is received correctly. In a communication between two (computing) entities, messages are rarely, if ever, sent without some type of control information. Figure 1.10 shows how a message, plus the control information, is sent in a frame or a packet. Both frames and packets must conform to a standard agreed upon by all entities in the communication.

What is transmitted as control information depends, in part, upon the protocol being used. At a minimum, the control information should signal to the receiver that a message is about to begin and that the message has ended. In situations where many entities could accept a message, the control information might contain a destination identifier, which identifies the intended destination entity. A source identifier is usually included to permit the destination entity to determine the transmitter of the message.

As an analogy, the frame can be considered an envelope into which a letter (the message) is placed. The destination is the address on the front of the envelope and the source is the return address.

Some protocols may require that responses be generated at certain intervals during a communication. These responses, which are also a form of control information, must be distinguishable from the information within the frame; otherwise, it will be impossible for the entity receiving the frame to determine whether it is a message or control information.

1.3.3
Error Checking and Recovery

During a communication, an entity may not receive the incoming message correctly, making it necessary for the receiving entity to indicate that fact to the transmitting entity. This is analogous to a conversation between two people; if the listener does not understand what is being said or missed what was being said because of some form of interference (such as a loud noise), the speaker may be asked to repeat what has just been said.

When a message is damaged or lost because of a fault on the communi-

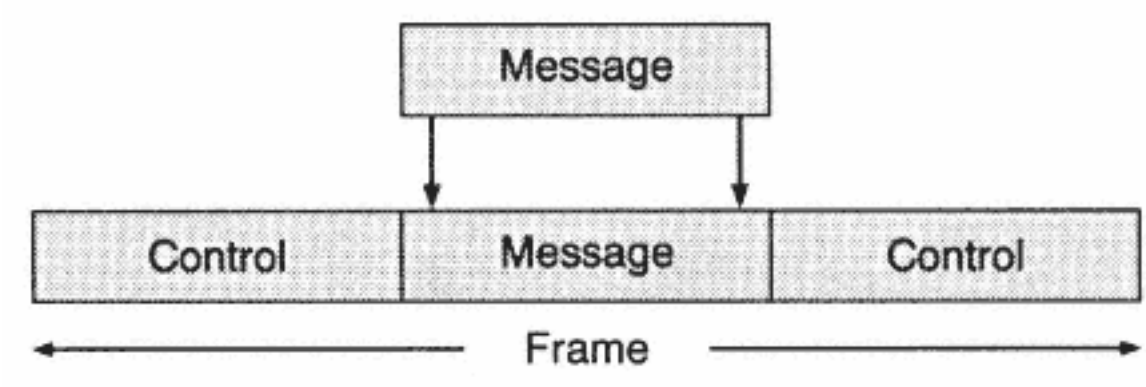

Figure 1.10
The relationship between a message and its control information.

cation channel, an error is said to have occurred. If a communication is to be reliable, the entities involved in it should be checking each message for errors and, if they detect one, have an agreed-upon procedure whereby the damaged or lost message can be recovered.

Errors can occur on the communication channel connecting computers through any one of a variety of faults, including

Signal attenuation. As a signal is transmitted through a medium, it must overcome resistances in the medium and in so doing may become weakened to the point where it is no longer possible to determine its meaning. For example, consider the problems involved in trying to determine what is being said by someone talking in whispers who is standing several hundred meters away. Signal attenuation can be overcome by either boosting the initial signal strength (i.e., getting the person to talk louder) or including amplifiers in the communication channel. There is a trade-off here: The signal cannot be boosted to the point where it is unintelligible or distorted to the receiver (consider the effect of yelling in someone's ear).

Noise. Changes in signal strength due to external interference or anomalies within the communication medium are known as noise. Categories of noise include:

White (or thermal) noise is background interference caused by the motion of electrons in the communication medium due to temperature. Unless the communication can take place at absolute zero, such interference cannot be eliminated.

Intermodulation noise occurs when a number of communications share the same channel (as in a full-duplex communication, described in Section 1.2) and one frequency interferes with another. In Figure 1.11, a station's transmission hardware is faulty, producing signals at the wrong frequency and thus causing intermodulation noise. In this case, station *B* is to transmit signals at frequency *Y*. but instead signals are produced at frequency *X*. thereby interfering with the transmission from station *A*. which is also transmitting at frequency *X*.

Signals that are sent electrically across wires can be subject to electromagnetic interference or impulse noise caused by events such as lightning or power surges.

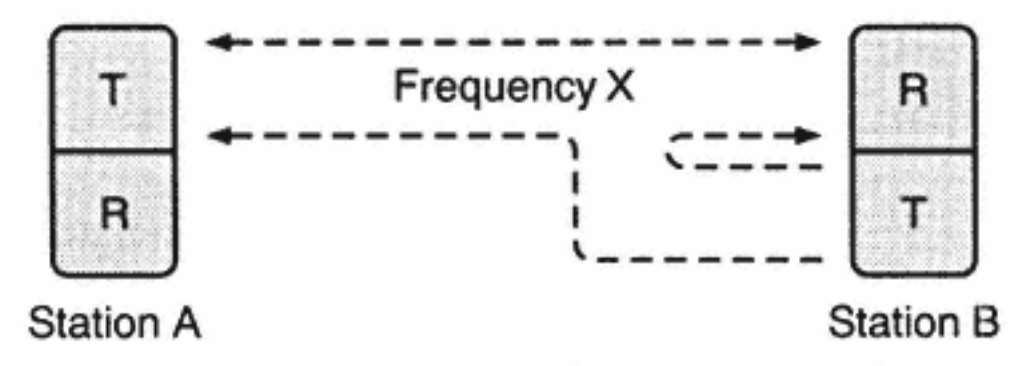

Figure 1.11
Intermodulation noise causing a frequency shift.

In addition to the above, a communication can be curtailed if the transmission medium is physically destroyed (say, in a fire).

The recovery from an error typically involves the retransmission of the message; that is, the transmitter begins retransmitting from the first message that the receiver missed. In many protocols, there is a limit to the number of times the transmitter will retransmit the same message. Should this limit be exceeded, the transmission is terminated since the flow of information approaches 0 as the entire communication is taken up with retransmissions of the same message and the receiver is not obtaining any useful information.

Similarly, recovery from an error may make it necessary to have mechanisms that ensure that both entities can remain in step with each other and that the transmitter does not flood the receiver with too much information. For example, during a conversation the speaker (the transmitter) often watches the listener for cues to indicate that what has been said is understood. If talking too fast, the speaker may be requested by the listener to slow down. This is known as flow control.

Computers can use flow control to limit the amount of information flowing across a channel. Simple terminal-host communications often use the control characters X-OFF and X-ON for flow control. Thus, typing X-OFF (CTRL-S) will cause the transmitting host to stop sending, thereby allowing what is on the terminal's screen to be read, and typing X-ON (CTRL-Q) will restart the output. More advanced protocols use various control sequences to achieve the same effect; these protocols will be discussed in subsequent chapters.

Frames (or packets) sent from a source to a destination can be lost or damaged, either through faults on the communication channel or by the destination entity losing frames because of processing overheads. The fact that frames can be lost or damaged means that the protocol should support mechanisms that permit the destination to determine if the incoming frame is in error or out of sequence (implying that a frame has been lost). These mechanisms include the following:

Out-of-sequence frames can be identified if each frame is sent with additional control information known as a sequence number. If each frame is sent with a unique number and both the source and destination agree upon the numbering scheme (as an example, 0, 1, 2, and so on), then if frames 0, 1, and 2 are transmitted but only 0 and 2 are received, the destination "knows" that frame 1 must have been lost.

Both the source and the destination must use the same algorithms for determining the next sequence number and agree upon the layout of the control information within the frame.

Once an error has been detected by the destination, it is standard practice to attempt to recover from it. In Figure 1.12, a typical recovery is shown

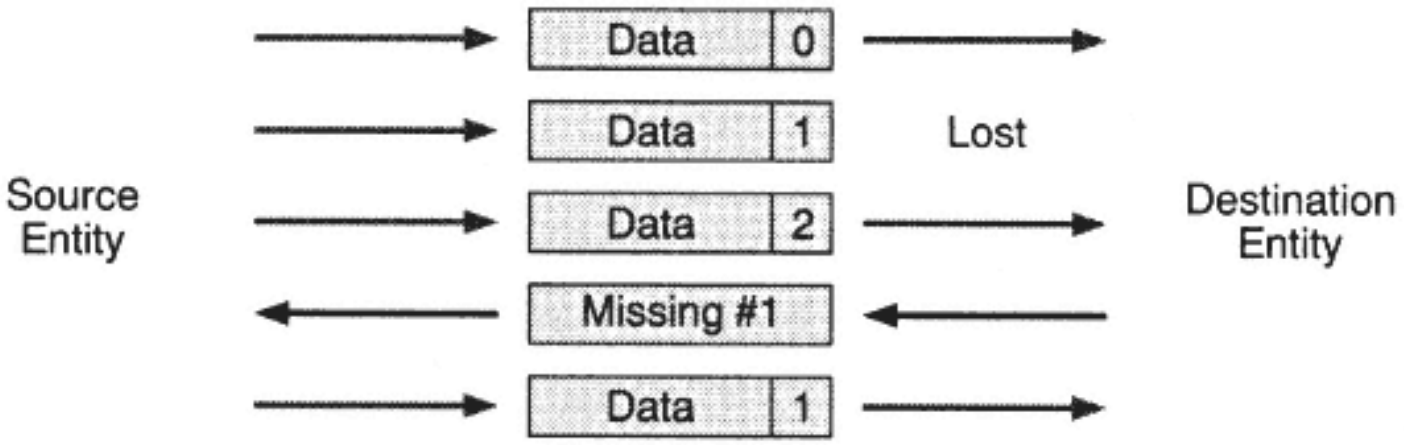

Figure 1.12
Error recovery using sequence numbers.

that entails having the destination send some form of error indication to the source, which may then retransmit the message.

Although the mechanism for recovery may seem simple enough, other problems can occur, including how the protocol handles the situation in which the error indication message is lost, and what procedures should be taken if the message sent from the source is always received in error.

Techniques for overcoming these and other problems will be discussed in later chapters.

1.3.4
Channel Utilization

In a communication, the information flowing between the entities must utilize some type of communication channel. For example, when two people talk on the telephone, the channel is the telephone system. Ideally, communication channels should be used as efficiently as possible, maximizing the amount of information passing from one entity to the other (the amount of useful information that flows between the transmitter and the receiver is known as throughput).

As an example of channel utilization (or throughput), consider the delivery of letters to someone's home in which the channel is the mail carrier. The carrier could deliver one letter at a time, rushing back and forth between the post office and the housea clear example of poor channel utilization and throughput. On the other hand, in an example of efficient channel utilization, the carrier delivers all the letters to the house at once, thereby minimizing the number of trips required and hence maximizing throughput.

Throughput can be maximized in a number of ways. For example, in a frame the ratio of control information to message should be as low as possible (that is, there should be fewer control information bytes than message bytes), otherwise, the channel is carrying more control information than useful information. Throughput can also be enhanced by increasing the speed of the channelthat is, the number of bits (or bytes) that are sent in a given period of time.

A third technique involves *how* the information in a frame is sent be-

tween the communicating entities. Typically, information is sent between computers as a stream of bytes (for example, frames are constructed out of a series of bytes). The channel between the two computers can be either serial, which means that one bit at a time is sent, or parallel, meaning that all eight bits making up the byte are sent at once; this is shown in Figure 1.13.

The obvious advantage of sending bytes in parallel is that it is eight times quicker than sending them serially. The trade-off is the cost of setting up a communication channel consisting of eight paths for parallel since eight times as many connections and paths are needed. For this reason, one usually finds parallel connections *within* the computer or between closely coupled computers; however, over distances of more than a few meters, most communications are serial.

Within the computer, bytes are moved in parallel. Before a byte can be transmitted serially, it must be converted from parallel using a hardware parallel-to-serial converter. Similarly, the stream of serial bits must be assembled into a (parallel) byte using a serial-to-parallel converter.

1.3.5
Synchronization and Timing

In any communication, certain steps must occur before others. For example, before a telephone conversation can take place, one person must dial the number of the other person, and before the person being called can answer, the telephone must ring. Once the called person answers the telephone, the conversation can proceed.

Protocols often require the entities involved in a communication to be in step, or synchronized. That is, one event is followed by another. Continuing with the telephone analogy, once the call has been established, synchronization is maintained, typically with one person making a series of statements to which the other responds. At the end of the conversation, both parties agree that the conversation is over and hang up.

Synchronization in a computer can occur at any number of levels. Two communicating processes must exchange synchronization information (such

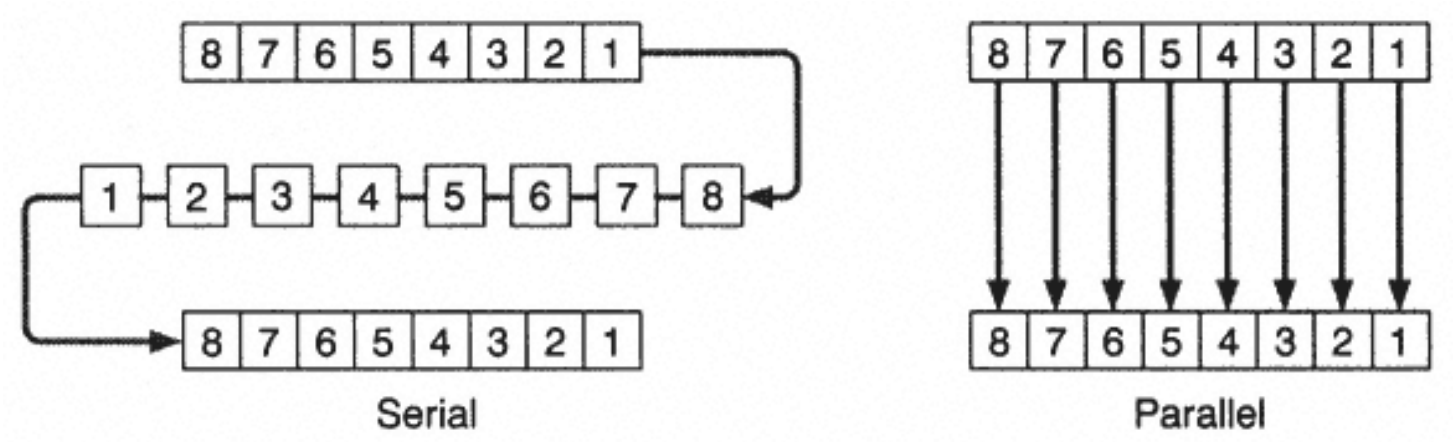

Figure 1.13
Serial versus parallel communications.

as sequence numbers, as discussed in Section 1.3.3) to ensure that no messages are lost. Synchronization is achieved by the use of an external timing device, typically a clock.

At lower layers, synchronization must be maintained by the communicating computers at the frame, byte, and bit levels. Since the computers are operating independently, the receiver must "know" when the next bit is to arrive. This can be achieved by having each computer maintain a clock that is in sync with the other computer's clock; at each tick of its clock, the transmitter sends a bit, and at each tick of the receiver's clock, a bit is read from the channel.

The term synchronous communication is applied to communication systems that can send a stream of bytes with only periodic timing information (typically attached to the start of the frame); the term asynchronous communication is applied to communication systems that require each transmitted byte to be prefixed with timing information.

Communicating entities can also use timing to ensure that the communication is still in progress. If an entity does not receive a response to a message within a given period, it can take action to determine why. The inactivity could have a variety of reasons: there may be a break in the communication channel (giving the false impression that the other computer is inactive), or the other computer may have crashed (thereby halting any communication activities on the channel). Many protocols require special frames to be sent periodically to allow each entity involved in the communication to determine that the others are still active.

Examples of synchronization and timing will appear throughout the remainder of the book.

1.3.6
Transparency

In any communication, the underlying facilities supporting it should be transparent to the entities involved. For example, the people involved in a telephone communication should not be concerned with the details of how the system actually places the call. The same argument can be applied to entities within a computer communication system, for similar reasons:

Application programmers should not be concerned with the hows and whys of the underlying communication system, as long as the facilities exist to get information from one machine to another. For example, the channel could be serial (or parallel) or the bytes could be sent synchronously or asynchronously, but these issues should be transparent to both the programmer and the application.

If an application is written for a specific communication system, it can become very difficult to transfer the software to another computer using a different communication system. Therefore, by layering the software and

making the underlying layers transparent to the application, it should be possible to transfer the software to another computer as long as the interface between the application and the lower layers is the same for both systems.

The term *transparency* can also be applied to the type of information transmitted. A transparent frame is one that can contain all possible characters (both control and non-control) as part of its message, while a non transparent frame is one that can contain only non-control (i.e., printable) characters.

1.4
Data Terminal Equipment

In its simplest form, a computer can be considered a device that when given input, processes it and produces output. If one computer is to communicate with others, it also requires some type of connection. A computer that serves as a point of collection, use, or transfer of data is referred to as Data Terminal Equipment or, more commonly, as a DTE, and is represented diagrammatically in Figure 1.14. Examples include ATMs (Automatic Teller Machines), point-of-sale terminals, and laser printers.

1.4.1
Input

To the vast majority of people who use computers, input takes place through the keyboard. Most keyboards are equipped with more than the standard typewriter "QWERTY" keys, including, for example, function and cursor control keys as well.

Over the past decade, a variety of specialized input tools have been developed, including optical scanners (to interpret the bar-codes associated with retail items such as supermarket goods), voice-to-digital units (to allow people to talk to the computer rather than type), magnetic strip decoders (to read the magnetic material found on the back of most charge cards), and pointing tools such as the mouse and trackball (to allow a user to select an object on the screen).

Although most DTEs support either ASCII or EBCDIC internally, input to the DTE need not be either. The PC keyboard is an example of an input device that produces its own character code, which must be converted to

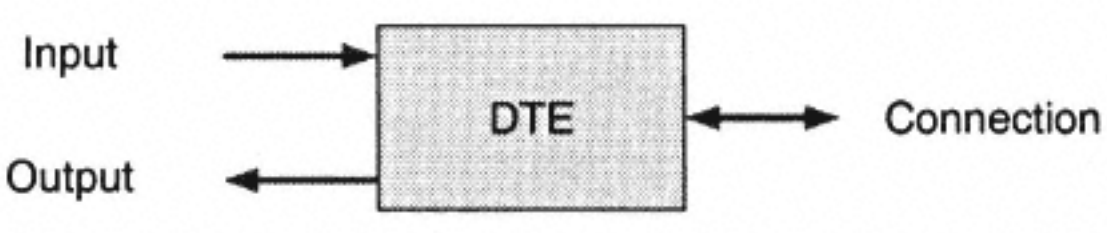

Figure 1.14
A DTE.

another character code (typically ASCII) before being processed. Similarly, a voice input device may produce 8-bit data values, but an individual byte probably has no direct equivalent in either ASCII or EBCDIC.

1.4.2
Output

For the most part, DTE output is visualoriginally textual either as hard (paper) or soft (cathode ray tube or CRT) copy. Today, output can be bit mapped images, including icons (still pictures on the screen), facsimiles (fax documents), and animation. Monochrome CRTs are being superseded by color and liquid crystal display (LCD) terminals.

A DTE need not be associated with input. For example, a laser printer (a DTE) produces output and is connected to a computer, but it need not support input.

Similarly, not all output need be printable text. Consider the following examples:

Both ASCII and EBCDIC support control characters that can be used to control the output text stream. Certain escape sequences, for example, are intended as control signals to the DTE rather than to the user.

Some DTEs support voice output for tasks that require listening rather than reading. For example, voice output is necessary for people with sight disabilities.

Output can be an *action* such as unlocking a door after a person has entered the password associated with the lock. In other words, the DTE's output is used to control an external device.

1.4.3
Connections

By definition, a DTE must have at least one connection to another DTE, otherwise communication cannot be achieved. Connections are typically physical, consisting of insulated wires or glass (optical fiber cable). In longer-distance communications, the DTEs need not be connected physically; information can be passed through the air (microwave communications) or through outer space (satellite communications).

It is worth noting that some DTEs do not support direct input and output. Some highly specialized computers, for example, do not communicate with humans; instead, all input and output passes through an external processor, which provides the link between the machine and its users.

1.4.4
More DTE Terminology

Probably the most common DTE interaction that users experience involves a terminal connected to a remote host (both the terminal and the host are examples of DTEs). Terminals usually operate in one of two modes:

Block mode. A block-mode terminal is one in which the information entered by the user is not sent to the remote DTE until the terminal receives an explicit signal from the user. Examples include airline reservation systems in which information is accepted from a client and verified before being sent to the remote host for processing. Block-mode terminals are typically CRTs.

Character mode. Character-mode terminals accept single characters entered by the user and forward them to the remote host. In some situations, the terminal will buffer one line of characters before sending it.

The characters entered by a user can be displayed from either the local or the remote DTE:

A remote echo occurs when a character is entered at the user's local DTE and the remote DTE is responsible for outputting it. Figure 1.15 illustrates the cycle of a remote echo: The local DTE accepts the character and transmits it to the remote DTE, which then performs whatever processing is required on the character and returns it to the local DTE for output. Remote echoing is often used in applications such as screen editors that do not expect the user to end each command with a delimiter such as a carriage return. Another term for remote echoing is echoplexing.

A local echo occurs when the user's local DTE is responsible for echoing each character entered by the user. The local DTE must still forward each character to the remote DTE (data sent from the remote DTE is displayed on the user's local DTE), as shown in Figure 1.16. Local echoing is often used by block-mode terminals; when the user types Enter, the information is transmitted to the remote DTE.

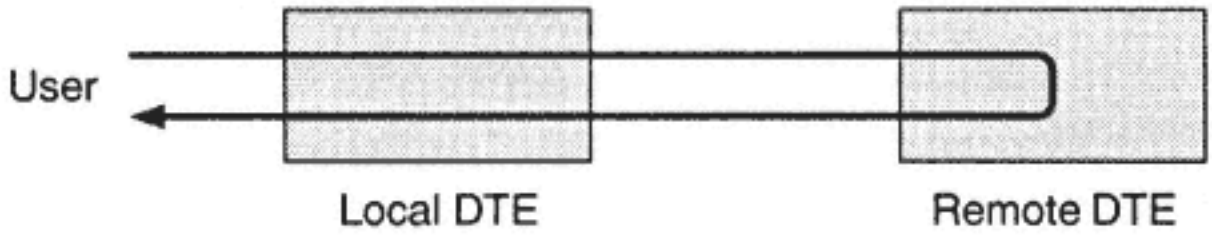

Figure 1.15
Remote echoing.

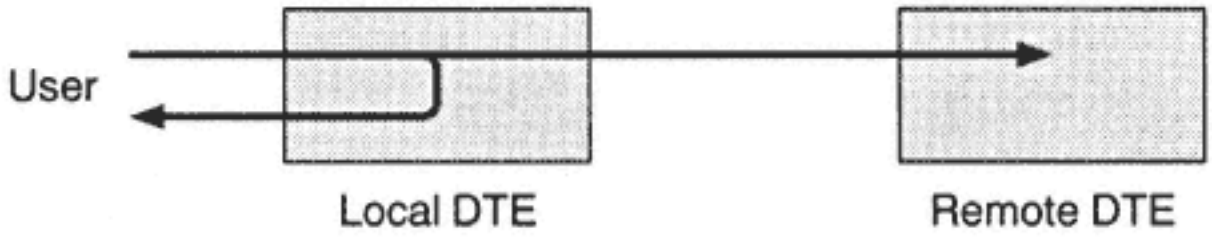

Figure 1.16
Local echoing (the local DTE echos each character before transmission).

1.5 Summary

This chapter introduced some of the terminology and concepts associated with data communications. Two points should be remembered:

1. Almost all aspects of data communications have a human parallel; when stuck trying to unravel a problem in data communications, think of it in human terms.

2. The only reason society functions is that it has rules; the same is true with data communications. Protocol standards exist to be adhered to; otherwise, there is no guarantee that the implementation or any applications that use it will succeed.

1.6 Exercises

Note that none of the exercises in this chapter refer to Commkit or the Commkit software. Exercises labeled "(A)" are simple questions, exercises labeled "(B)" refer to commands issued to an operating system, and exercises labeled "(C)" require some software design and implementation.

1. (A) Describe the medium and suggest the type of channel (either simplex or duplex) for the following communications:

    a. A person talking.

    b. A pigeon carrying a message.

    c. Television.

2. (A) Which of the following communications could be described as halfduplex? As full-duplex?

    a. Two polite people talking.

    b. Two angry people yelling.

    c. A message sent by smoke signals.

Of those that are half-duplex, what is the line turnaround signal (if any)?

3. (A) Assuming that most, if not all, human speech falls into the "halfduplex" category, identify some gestures and phrases that people use to indicate that they have finished talking (i.e., the point at which the line can be turned around).

4. (A) Describe the steps required to convert a byte stored in parallel to serial. Consider the hardware required for such an operation.

5. (A) Repeat the previous question, but convert the byte from serial to parallel.

6. (B) Experiment with several computers to determine if they support X-OFF (CTRL-S) and X-ON (CTRL-Q). What happens if CTRL-S is typed and then a command and finally CTRL-Q?

7. (C) Examine several computers to determine their character codes (ASCII or EBCDIC). Try implementing the following C program to display the computer's different characters:

```
#include <stdio.h>
main()
{
unsigned char ch;
for(ch = 0; ch < 256; ch++)
     printf("%d %c\n", ch, ch);
}
```

Now, explain why the above program runs forever, printing the characters and their numeric equivalents.

8. (C) In certain situations a file containing ASCII characters (such as a program or electronic mail) may be sent to a machine that supports only EBCDIC characters. Similarly, EBCDIC files may be sent to ASCII machines. When the file arrives at the destination, it is necessary to convert the file contents to the character code of the destination machine.

Write a conversion program that will take a file in either ASCII or EBCDIC and convert it to the other character set. Ideally, the program will operate as follows:

```
C:\> convert DIR in-file out-file
```

The option `DIR` should indicate the direction (either `EA`EBCDIC to ASCII, or `AE`ASCII to EBCDIC).

To examine the contents of an EBCDIC out-file, use the system's dump utility. Create a dummy EBCDIC file by generating a file of all 256 possible byte values. If a character does not map directly into an ASCII (or EBCDIC) character, replace the character with a question mark.

9. (A) A DTE displays every character typed by a user *twice*. Explain what is happening and how to correct it.

10. (A) A DTE doesn't display anything that is typed, but the remote DTE responds when a carriage return is pressed. Explain what is happening and how to correct it.

2
Commkit

2.1
Introduction

Before a computer can communicate with a user or another computer, it requires hardware to support the communication and software to control the communication. This chapter examines some of the hardware supplied with a typical PC and considers the software required to control it. The chapter then introduces Commkit, a software tool that supports the development of practical data communications software. An example of an application that uses some of the PC's hardware and Commkit is also presented.

2.2
The 8086 Family and Its Hardware

The 8086 is the generic name given to a family of microprocessors developed by Intel and supplied with personal computers such as the IBM PC and its clones. At present, there are some six members of this family that can support Commkit:

8086the "original" 8086 processor (there were earlier versions, such as the 8080 and 8008, but they are not of concern here).
80186an enhanced version of the 8086, with several new instructions.
80286a faster, more powerful version of the 8086 with enhancements for segmented memory management and multitasking.
80386a still more powerful version of the 80286, supporting segmented and paged memory management as well as additional instructions for 32-bit arithmetic and bit manipulation.
80486a faster version of the 80386 with built-in floating-point hardware.
Pentiumoriginally called the 80586, a still faster version of the 80486, that uses features such as *pipelining* to execute instructions in parallel.

For the most part, all members of the 8086 family are upwardly compatible, meaning that software developed for an earlier version of the processor should be able to run on a later one.

Although recent versions of the 8086 family (such as the 80486) are more powerful than earlier ones (such as the 8086), the mechanisms whereby the external hardware (or *devices*) is accessed have essentially remained unchanged, once again to allow the upward compatibility of software.

2.2.1
8086 Devices

Information is supplied to and received from an 8086 processor through devices such as the keyboard, screen, disk, light-pen, serial port, and mouse. Although there seems to be a limitless supply of possible devices that can be attached to a PC, the standard PC configuration generally allows only eight (typically, the clock, keyboard, screen, hard disk, floppy disk, printer, and two serial communication interfaces). AT configurations allow a total of 15 devices.

Devices are not accessed directly (as, for example, memory is); instead, an 8086 processor accesses a device through one or more ports. To ensure that the PC can distinguish between them, all devices are assigned one or more unique port numbers. Although only a limited number of devices can be attached at any one time, there are some 216 possible ports available on the PC. The number of ports associated with a device depends, in part, upon the number of functions it performs. For example, the clock is associated with four ports; the keyboard uses two.

Ports can be accessed through software using two "low-level" instructions. The `in` instruction allows a port to be read, while the `out` instruction allows it to be written. The `in` instruction accepts a port number and returns the value (a byte) associated with that port; the `out` instruction requires both a port number and the byte to be written to the device. The "high-level" Turbo C counterparts of these instructions are

`unsigned char inportb (int portid);`returns an 8-bit byte associated with port `portid`.
`void outportb(int portid, unsigned char value);`writes an 8-bit quantity, `value`, to the port `portid`.
`void outport(int portid, int value);`writes a 16-bit word, `value`, to the port `portid` and port `portid + 1`.

The types of operation (i.e., reading or writing) that can be performed on a port depend upon the functions of the device that the port supports. Some ports, such as the input buffer associated with the keyboard (port number `0x60`), are for reading, while others, such as the clock command register (port number `0x43`), used for programming the clock, are for writing. Finally,

some ports can be both written to and read from. One example, the keyboard status register (port number `0x61`), can be read (to obtain the status of the keyboard) and written to (to signal the keyboard that the supplied character has been accepted).

Of the eight "standard" devices that can be supported by a PC, only two are of direct interest at this moment, the keyboard and clock, since they are both used by the Commkit software.

The Keyboard

The keyboard is an input device that allows a user to supply information to the processor in alpha-numeric format. The number of keys on the keyboard and their layout depend upon the type of PC and the keyboard's manufacturer.

Although the 8086 processor uses the ASCII character code, the PC keyboard does *not* generate ASCII characters; instead, each character on the keyboard is associated with a one-byte scan code. The scan code is returned to the processor via port `0x60`.

IBM has defined a scan code for each key (to remain compatible with the IBM PC; keyboards built by other manufacturers must generate scan codes that correspond to those found on the IBM PC, regardless of where the keys are placed on the keyboard). For example, the Escape key (ESC) generates scan code `0x01` and the Delete key (DEL) generates scan code `0x53`. Since the value of the scan code does not correspond to a specific character code (such as ASCII or EBCDIC), the scan code must be translated into the character code required by the application. The translation is done through software within the processor, mapping the scan code into the equivalent character code character, typically with a translation (or mapping) table as shown in Figure 2.1.

The processor is signaled *twice* whenever a key is touched: once when the key is pressed (generating the make scan code), and once when the key is released (generating the break scan code). Regardless of whether the scan code is a make or a break, its lower seven bits identify the key. The eighth bit is cleared by the keyboard to indicate a make and is set to signal a break.

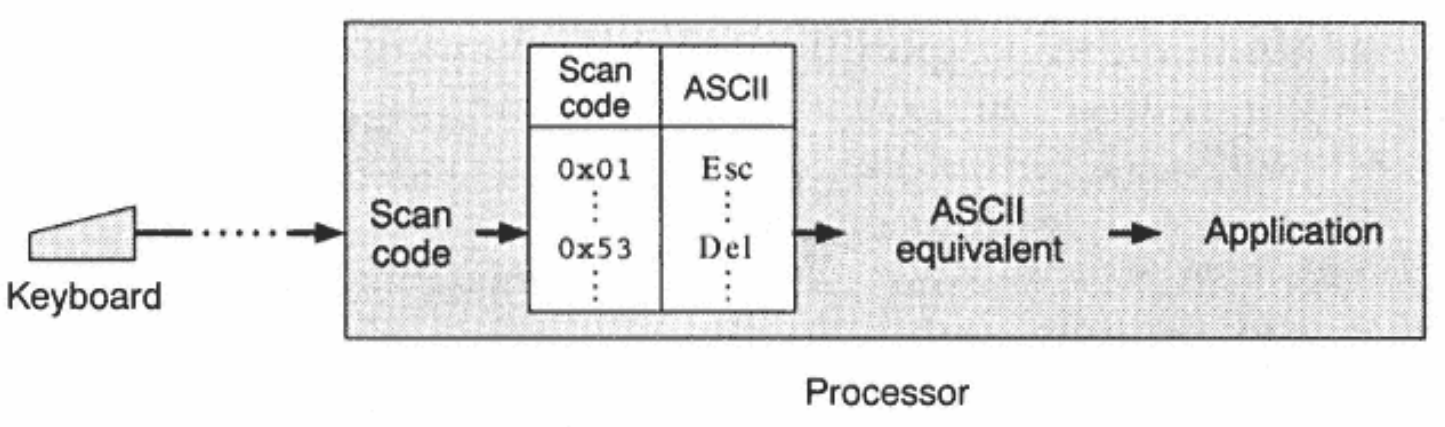

Figure 2.1
Steps in converting a keyboard scan code into ASCII.

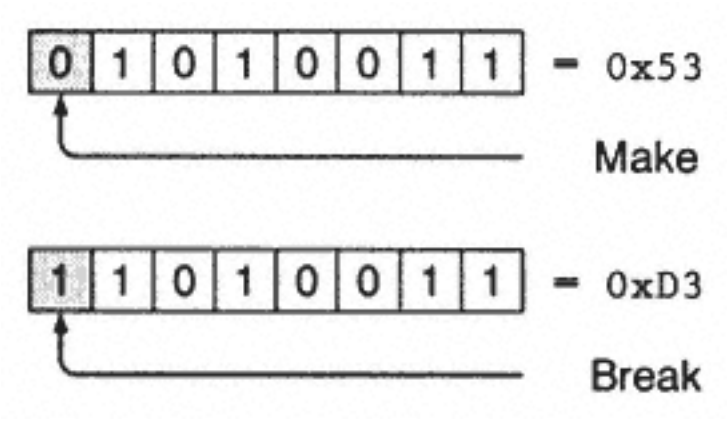

Figure 2.2
The make and break scan codes for DEL.

In Figure 2.2, the make scan code for the Delete key (generated when the DEL key is pressed) is `0x53`, while the break scan code (generated when the DEL key is released) is `0xD3`.

The keyboard can generate 127 different scan codes (with values 1 through 127). Scan code `0x00` is reserved to allow the keyboard to expand to more than 127 characters. Keys outside the 127-character range are sent as *two* scan codes, `0x00` and the scan code of the key. If the keyboard is equipped with an "auto-repeat" feature (that is, when a key is held down, the keyboard eventually starts to repeat the character), it sends each character as a make, then as a break, eliminating the need for the user to continuously press and release the same key. The keyboard is also equipped with a 16-byte buffer for storing scan codes until they are read by the processor.

The keyboard software obtains the scan code by reading port `0x60`. It must then signal the keyboard that the character has been reada seemingly convoluted process involving the following steps:

1. Read port `0x61` to obtain the keyboard status.

2. Write port `0x61` with the keyboard status or'ed with `0x80`.

3. Write port `0x61` with the original keyboard status.

The "flip" of the high-order keyboard status bit signals that the scan code has been read, thereby allowing the keyboard to remove the scan code from the keyboard buffer. An algorithm that converts all scan codes into a single case (i.e., all upper or all lower) can be implemented by ignoring any incoming make scan codes and processing the break scan codes only:

1. Obtain the scan code from port `0x60`.

2. Signal the keyboard via port `0x61` that the scan code has been obtained.

3. If the scan code is a break:

    a. Convert it into a character, ideally using a mapping table that consists of the ASCII characters associated with the different scan codes. The scan code is used as a subscript into the mapping table.

    b. Process the character.

The above algorithm is limited in that it supports a single case and doesn't permit control characters. Fortunately, the make/break cycle allows a program to determine which keys have been pressed and in what order, thereby

permitting the keyboard software to distinguish between sequences such as A-SHIFT, SHIFT-A, or even SHIFT (release) A. For example, the "A" key on the keyboard can generate one of a number of ASCII values, depending upon the other keys that were pressed beforehand and how the software interprets the sequence of keys being pressed (note that the keyboard distinguishes between left and right SHIFTS, and that CTRL is an abbreviation of Control). Various key combinations are shown in Table 2.1.

If multiple keys are pressed (for example, to obtain a capital letter, a control character, or a special sequence such as CTRL-ALT-DEL), the keyboard software must maintain state information about these keys, since the scan codes are supplied to the processor one at a time. At a minimum, the keyboard software should be able to "remember" whether the CTRL, ALTMODE, LEFT SHIFT, and RIGHT SHIFT keys have been pressed. The state of any of these keys can be maintained as a Boolean condition (a key is pressed or it isn't), with the initial state of each key being `FALSE`. Whenever one of these keys is pressed (i.e., the make scan code), the state can change to `TRUE`, and when the key is released (i.e., the break scan code), the state can change to `FALSE`.

This also means that two mapping tables are required, one for unshifted characters and the other for shifted characters. There is not necessarily a one-to-one correspondence between the tables (for example, "a" to "A"), since some unshifted characters don't have a shifted equivalent (such as "1" and "!").

ASCII control characters are those less than `0x20` (space) and can be obtained by and'ing the character with `0x1F`. Similarly, ALTMODE characters are those greater than `0x7F` (DEL) and are generated by or'ing the character with `0x80`.

The Clock

The clock (or more correctly, the 8253 timer chip) is used to supply the PC with a regular, periodic clock pulse that can be used to control various actions.

Table 2.1 Keystroke interpretation.

First Key	Second Key	Third Key	Result	ASCII Value
A			a	`0x61`
SHIFT LEFT	A		A	`0x41`
SHIFT RIGHT	A		A	`0x41`
CTRL	A			`0x01`
ALTMODE	A		β	`0xE1`
ALTMODE	SHIFT	A	⊥	`0xC1`
CTRL	ALTMODE	A	ü	`0x81`

Table 2.2 Clock channels and functions.

Channel	Function
0	System timing and counting
1	Memory refresh (via DMA controller)
2	PC speaker (for sound effects)

The 8253's *three* independent timing channels that are used by the PC are shown in Table 2.2.

Of the three channels, channel 0 can be used by programs (such as MS-DOS and Commkit) to control hardware and software access to the PC. Channel 1 *must not* be changed, since this can result in the loss of the contents of the PC's memory. Channel 2 is not used by Commkit.

Internally, the 8253 has a 1.19318 MHz clock that supplies each timing channel with 1,193,180 clock pulses each second. Since most applications do not require this accuracy of timing, each channel is associated with a programmable 16-bit counter that can be decremented by the timer chip on each clock pulse. When the counter reaches 0, the application can be informed. As an example, assume that an application requires the clock to signal the processor 1000 times a second (in other words, once a millisecond). The counter must be initialized to a value that will reach 0 after one millisecond has passed. Dividing the clock speed (1,193,180) by 1000 gives 1193; setting the counter to 1193 results in the counter reaching 0 after approximately one millisecond.

The 8253 clock is associated with four ports. Ports `0x40`, `0x41`, and `0x42` are the clock counter registers and are used to supply the initial clock values to channels 0, 1, and 2, respectively. Port `0x43` (the clock command register) allows the programmer to specify how a clock is to be used as well as how it is to be initialized. For example, once a clock's counter reaches 0, the clock can be programmed to load itself with the original counter value (stored in the clock latch) and repeat the cycle. Alternately, it can be programmed to stop at 0 (this is known as single-shot mode). Similarly, the value loaded into the clock latch (through the clock counter register) can be the full 16 bits (obtained by writing the low-order and then the high-order byte to the clock counter register) or simply half of the clock value (i.e., either the low-order or the high-order byte).

2.2.2 Accessing Devices

Although a device can be accessed at any time (through one or more of its ports), it is not always advisable to do so. For example, the keyboard register can be read regardless of whether or not the user has typed a character.

Accessing a device before it is ready can result in the duplication of information (reading the serial communication interface more than once before a new character has arrived will result in a copy of the original character being returned) or the loss of information (writing to the serial communication interface before the last character has been sent can result in the new character overwriting the previous one). To avoid situations in which data is lost or duplicated, most devices are able to signal their status to the processor. Typically, the status indicates whether the device has information for the processor or is ready to accept more information.

The status of the device can be obtained either by the processor polling the device or by having the device interrupt the processor.

Device Polling

The state of a device can be obtained by reading one or more ports associated with the device. For example, it is possible to configure the clock so that it counts down to 0 and stops. By polling the port associated with the clock, a program can determine whether the clock has reached 0. Software for device polling is typically written as a loop known as a polling loop:

```
for(;;)
{
    if (devicel_ready()) service_device_l();
    if (device2_ready()) service_device_2();
}
```

There are, however, at least two drawbacks to using device polling: First, the processor performs no useful function other than polling; second, if a device happens to generate data faster than it takes the processor to execute the polling loop, data can be lost. Consider, for example, the following polling loop:

```
for(;;)
{
    if (device1_ready()) service_device_1();
    if (device2_ready()) service_device_2();
    if (device3_ready()) service_device_3();
    if (device4_ready()) service_device_4();
}
```

If `device1` supplies data faster than it takes the processor to check each device in the polling loop, there is a possibility that data from `device1` will be lost. A common trick to overcome this problem is to poll the fast device more than once in the polling loop.

Interrupts

Ideally, what is required is a mechanism to signal the processor only when a device needs to be serviced or has information to supply to the processor. This allows the processor to perform tasks other than device polling (for example, a user can type information at a keyboard while other information is being written to a disk). Most processors, including those in the 8086 family, allow devices to signal, or interrupt, the processor when a condition has been reached, in this way overcoming the limitations associated with device polling. For example, instead of the software polling a disk to determine if a block of data has been written, the disk itself can inform the software that the data has been written.

When a device interrupts the processor, a number of things occur. First, the task currently being run is suspended while the processor handles the interrupt. Second, a procedure known as an interrupt handler (or interrupt service routine) must be activated. The interrupt handler is responsible for servicing the interrupt (that is, determining why the interrupt has occurred and what to do about it). Third, the suspended task must be resumed once the interrupt handler is finished.

When an interrupt occurs in an 8086 processor, the task currently executing is suspended by pushing the instruction counter and the status flag on the stack, thereby permitting control to be returned to the task once the interrupt has been serviced. To ensure that no further interrupts occur during the handling of the first interrupt, the 8086 processor disables them (i.e., if other interrupts occur, they are blocked until the processor either explicitly enables them or resumes execution of the interrupted task).

Each device is associated with a unique interrupt number that the processor obtains when the interrupt occurs. This number is used as an index into the list of interrupt vectors stored in segment 0. Table 2.3 shows the

Table 2.3 "Standard" PC Interrupt Vectors as assigned by DOS.

Interrupt Number	Interrupt Vector Location	Device
0x08	0x20 - 0x23	Clock
0x09	0x24 - 0x27	Keyboard
0x0A	0x28 - 0x2B	From slave 8259
0x0B	0x2C - 0x2F	Serial port
0x0C	0x30 - 0x33	Serial port
0x0D	0x34 - 0x37	Hard disk
0x0E	0x38 - 0x3B	Floppy disk
0x0F	0x3C - 0x3F	Printer

assignment of interrupt numbers, interrupt vector locations, and the eight "standard" devices (note that there is no obvious relationship between the orts associated with a device and the device's interrupt number).

When a device causes an interrupt, the instruction counter is assigned the value of the device's interrupt vector. Control is then passed to the interrupt handler. Once the interrupt has been serviced, the stack is popped, restoring the original task's instruction counter and status flag. Interrupts are reenabled because the interrupt enable bit is set in the status flag.

Although the 8086 processor is designed to handle up to eight external devices, only a single interrupt line connects the processor to the outside world. This means that without some form of additional hardware, at most one external device can be connected to the processor. Fortunately, hardware such as the Intel 8259 Interrupt Controller has been designed to share the single interrupt line between eight devices. Figure 2.3 shows the relationship between the devices, the 8259, and the 8086 processor.

The figure shows that instead of interrupting the processor directly, a device first signals the 8259, which then interrupts the 8086 processor using the single interrupt line. The processor determines which device is interrupting by obtaining the device's number from the 8259. The 8086 processor uses the device number to access the list of interrupt vectors that indicators which interrupt handler should be activated.

The 8259 permits the programmer to select the devices that are to interrupt the 8086 processor by writing a 1-byte interrupt mask to the 8259. Each bit in the mask corresponds to one of the eight devices. As shown in Figure 2.4, device priority is indicated from right to left, with the clock having the highest priority and the printer having the lowest.

A bit value of 1 in the interrupt mask indicates that any interrupts coming from the device are to be ignored, and a bit value of zero means that the device is allowed to interrupt the 8086. For example, to permit clock, keyboard, and printer interrupts, the interrupt mask would be set to `0x7C`, as shown in Figure 2.5.

The 8259 interrupt mask is accessed through port `0x21`, the interrupt mask register. The example from Figure 2.5 could be implemented as follows:

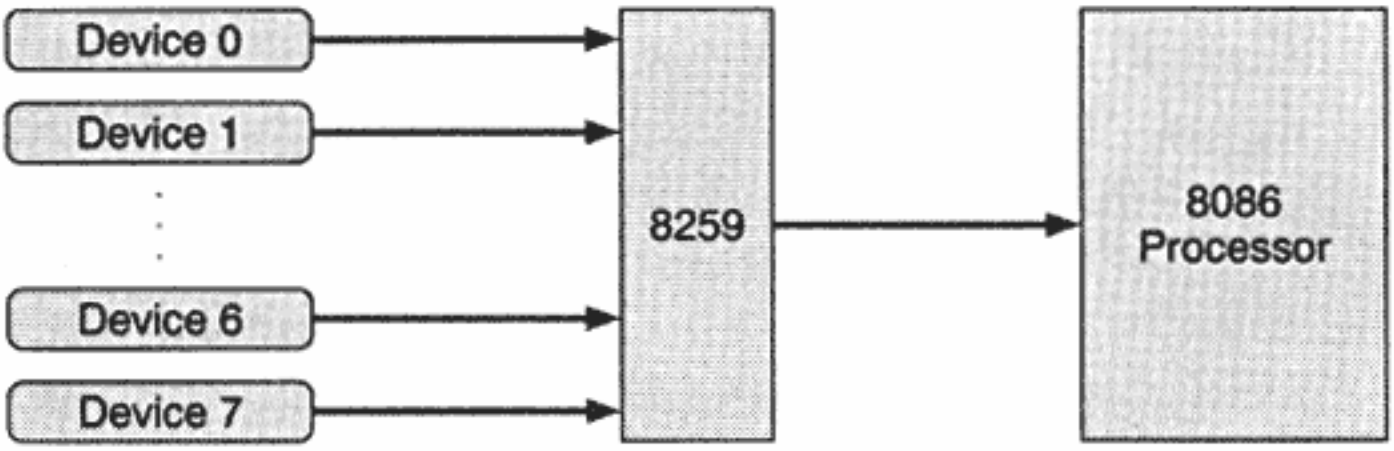

Figure 2.3
Connecting the 8259 to devices and an 8086 family processor.

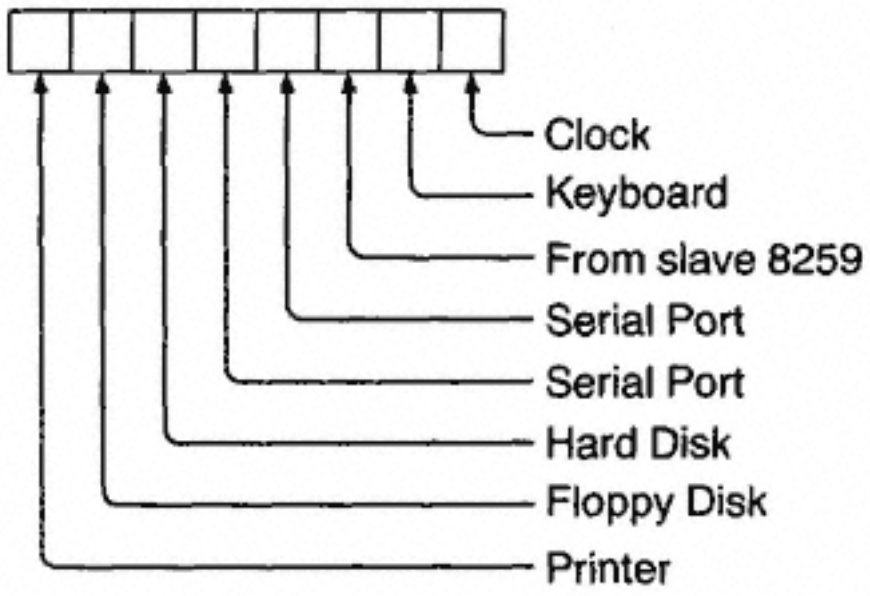

Figure 2.4
The 8259 interrupt mask.

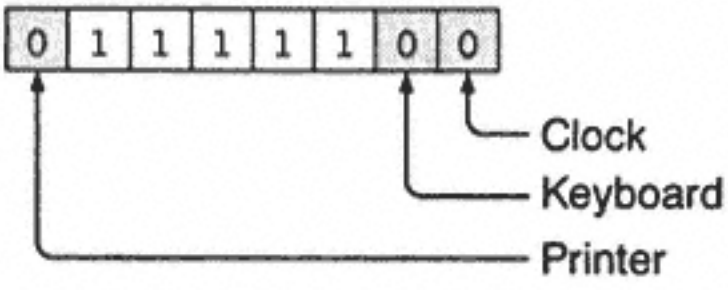

Figure 2.5
Interrupt mask allowing clock,
keyboard, and printer interrupts.

```
#define INT_MASK    0x21
#define CLKENA      0xFE  /* Clock enable:
                             11111110 */
#define KEYENA      0xFD  /* Keyboard enable:
                             11111101 */
#define PRTENA      0x7F  /* Printer enable:
                             01111111 */
. . .
outportb(INT_MASK, CLKENA & KEYENA & PRTENA);
```

For each device selected, there must be a corresponding interrupt handler and the interrupt vector associated with the device must contain the entry point of that handler. Results will be unpredictable if either the interrupt handler is missing or the interrupt vector contains an invalid entry point, since control will be passed to a location that does not service the interrupt.

If several devices interrupt simultaneously, the 8259 signals the processor with the highest-priority interrupt. All other devices (with lower-priority interrupts) are kept waiting. The keyboard interrupt has the highest priority of all the other devices (except the clock) to ensure that special sequences such as CTRL-ALT-DEL are not blocked. The interrupt handler should be kept as short as possible, since *all* interrupts are blocked while the interrupt handler is active unless the interrupt handler explicitly enables them. An unduly long interrupt handler can result in interrupts being lost. Once the interrupt

has been serviced, the 8259 must be informed so that any pending (or any subsequent) interrupts can be signaled. This is done by writing `EOI` (end-of-interrupt, a value of `0x20`) to the 8259 Operation Command Word register (port number `0x20`).

2.3 The Commkit Software

By itself, the hardware described in the previous section performs no useful function without software controlling it. For example, simply pressing a key on the keyboard does not mean that the character will appear on the screen; software is required to read the scan code from the keyboard, convert it into a character, and then display the character on the screen.

Few computer users have the time or the patience to write their own software to control the underlying hardware; instead, they rely upon software written by other people who know the hardware. This software, sometimes referred to as an operating system, kernel, or monitor, is intended to hide the idiosyncracies of the hardware by offering a common, well-defined interface to the user. Examples of operating systems include MS-DOS, Windows, UNIX, and VMS.

In addition to hiding the hardware, many operating systems support the pseudo-concurrent execution of a number of tasks, referred to as processes. Depending upon the operating system, a process may be associated with all the actions invoked by a user (as in VMS, where a single process is used for all tasks), or each action invoked by the user may result in a new entity (as in UNIX, where the user's process spawns separate child processes to perform a specified task, such as editing a file or compiling a program). In operating systems that support multiple processes, it is often both necessary and useful to allow the transfer of information between processes (for example, the output of one process may be the input to another). Operating systems that allow processes to communicate are said to support interprocess communication, or IPC. Interprocess communication is an extremely useful feature that is also a cornerstone in many aspects of computing science, including data communications, distributed systems, and object-oriented programming.

One method of implementing interprocess communication is through the use of messages. A message is simply a data structure passed between the communicating processes. As in any other communication, all entities involved in interprocess communication must agree to a common protocol covering things such as the the format of the message and the required control information. The underlying operating system is responsible for ultimately controlling the communication and supplying the message to the intended destination process. Figure 2.6 shows a typical cycle involving a process (referred to as the client) sending a request message to a second process (the server); the second process replies with a response message to the client.

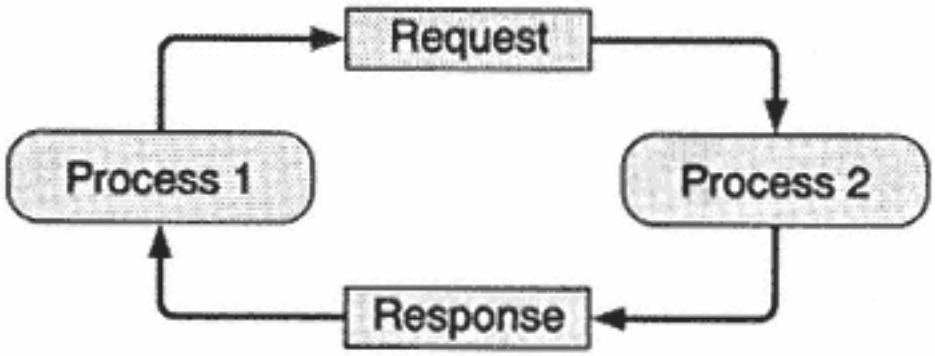

Figure 2.6
The requestresponse cycle.

Systems based upon the requestresponse message cycle are often referred to as client-servers.

Although MS-DOS supports neither concurrent execution of processes nor interprocess communication, it does run on the PC and many powerful tools have been written for it. To facilitate the teaching of data communications in an MS-DOS environment, it is necessary to either modify MS-DOS or develop a teaching tool that can be invoked by the user. It is possible to modify MS-DOS; however, without a detailed description of the MS-DOS internals, a course in data communications can degenerate into a painful exercise in debugging, with little time left over for communications. The approach to teaching data communications adopted in this book is to keep MS-DOS and its tools intact, but to use a stand-alone program known as Commkit.

Commkit is a small, event-driven monitor that supports multiple processes and message passing. A Commkit executable program is like any other MS-DOS executable program in that the program is invoked by typing its name after the DOS prompt. Once executing, Commkit controls the PC and allows the various processes to function and communicate. When it has finished executing, control is returned to MS-DOS.

2.3.1
Overview of Commkit

In Figure 2.7, the structure of a Commkit program is represented as a series of layers built atop the PC's hardware, each performing a specific function.

The Commkit software is supplied on a diskette containing about 90 files.

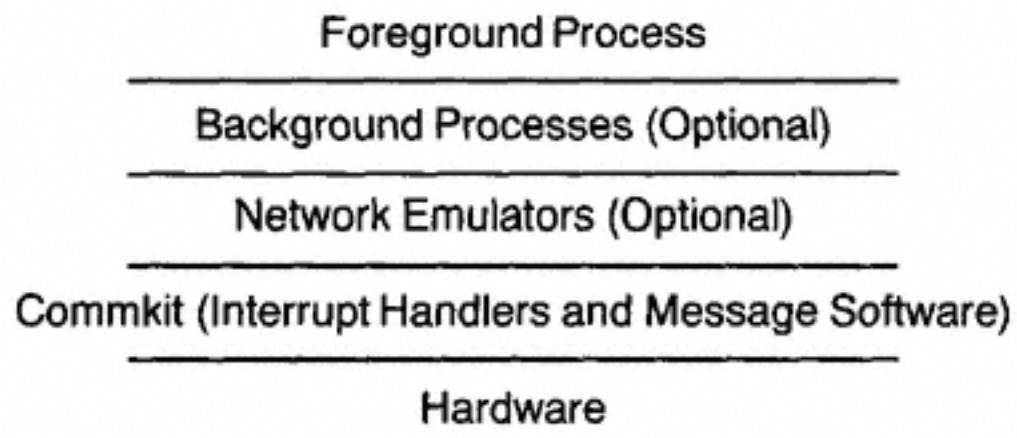

Figure 2.7
The structure of a Commkit program.

The various files consist of message-handling routines, network emulators, interrupt handlers, and a number of applications illustrating different aspects of data communications. All of the software is written in C and has been compiled and tested using both Turbo C and Turbo C++ (version 3.0).

A Commkit executable program is made by linking one or more modules containing a foreground process with the Commkit interrupt handlers and message-handling software. When necessary, modules containing additional processes (known as background processes) or those containing network emulators can be included to make an executable program.

All entities (i.e., foreground and background processes, network emulators, and interrupt handlers) communicate via messages using two communication primitives: `send()` and `recv()`. Each entity is associated with a unique process identifier, a message queue, an entry point known to Commkit, and a series of procedures implementing the process. For two entities to communicate, the source entity (i.e., the originator of the message) sends the message to the destination entity by calling `send()` with the destination entity's identifier and the message. The message remains on the destination entity's message queue until the destination entity accesses the message by calling `recv()`.

Every Commkit executable program requires a mainline procedure (i.e., the entry point from MS-DOS), called `main()`. Commkit must be initialized before it can be used. This is done in the mainline with a call to `initialize()` (defined in `commkit.c`) that sets up the message queues, interrupt handlers, and various support routines. Upon completion, control is returned to the mainline. Once initialized, any background processes are "registered" with Commkit using the `register_background` procedure (see Section 2.3.3). Finally, control is passed to the foreground process. Execution continues until control is returned from the foreground to the mainline, at which point `rtn_to_dos()` is called, restoring the interrupt vectors to what MS-DOS expects them to be.

Although the basic message-handling paradigm is true for all entities, there are minor variations that warrant an examination of each type of entity supported by Commkit.

2.3.2
The Foreground Process

The foreground process is simply an algorithm implemented as a C function that is called by the mainline. The Commkit software (consisting of the foreground process, any background processes, and the interrupt handlers) will continue to run until a condition is reached that causes the foreground process to return to the mainline, at which point Commkit is terminated and control is returned to MS-DOS.

The basic structure of the foreground process is a large loop that is terminated when some condition is reached (for example, the user typing

CTRL-C). Within the loop is a call to `recv()` that allows the process to receive messages from other entities:

```
void foreground_process()
{
int running = TRUE;
while (running)
{
     recv( /* recv() parameters */ );

     /* Some condition is reached. . . */
     running = FALSE;
}
}
```

The foreground process is interruptable by any of the interrupt handlers (interrupts are transparent to the foreground process). Control remains with the foreground process until it attempts to receive a message. If a message is available, it is returned immediately to the foreground process; otherwise, the process is blocked and any background processes with pending messages are allowed to execute. As soon as a message arrives for the foreground process, control (and the message) are returned to it.

The foreground process is associated with the process identifier `APPLICATION`; all messages sent with a destination of `APPLICATION` are placed on the foreground process's queue. There is only one foreground process.

2.3.3
Background Processes

A background process, like the foreground process, consists of an algorithm implemented as one or more C functions; however, background processes differ from the foreground process in a number of ways:

1. There can be any number of background processes (zero or more), whereas there can only be a single foreground process.

2. Background processes are executed only when the foreground process is idle (i.e., waiting for a message) and a message is available for the background process.

3. Background processes must be *registered* with Commkit before they can receive messages. Each background process is registered separately with a call to `register_background()` (defined in `general.h` and implemented in `commkit.c`):

```
void register_background(int proc_id, void (*proc_name)());
```

4. Background processes should not contain software that waits for multiple messages or implements infinite loops. Failure to observe this can result in Commkit losing all of its message buffers as messages go unread.

5. The entry point (i.e., a procedure) associated with a background process must be declared as type `void`.

Since control is only to remain with a background process for the time it takes to handle a single message, the following points must be considered when writing a background process:

If local variables are used, they are undefined upon entry to the process.

Static variables allow values to be kept while the process is inactive.

Avoid using global variables to store values between calls to a background process. Global variables can be changed by other routines; for example, since interrupts can occur at any time (i.e., clock or keyboard), it is possible that the value of a global variable can change while a background process is executing, potentially leading to unpredictable results.

Each background process is identified with a unique identifier; five such identifiers are defined in `general.h`: `BACKGROUND_1`, `BACKGROUND_2`, `BACKGROUND_3`, `BACKGROUND_4`, and `BACKGROUND_5` (more can be added if needed). However, if more are required, the global constant `TOTPROCS` (also defined in `general.h`) must be increased to reflect the total number of processes and interrupt handlers.

If control remains in a background process (for example, an infinite loop is entered), execution can be terminated by typing CTRL-ALT-DEL, which causes the keyboard interrupt handler (described below) to call `rtn_to_dos()`.

2.3.4 Interrupt Handlers

Commkit supports five external devices (keyboard, clock, two serial interfaces, and the parallel port interface), each of which is associated with an interrupt handler. The software associated with each interrupt handler can be found in `commkit.c`. An interrupt handler is associated with a process identifier, which means that it can both send and receive messages. The process identifiers are listed in Table 2.4 and defined in `general.h`.

Of the five external devices supported by Commkit, only two are considered at present: the keyboard and the clock. However, before examining the interrupt handlers, the available support software is discussed.

Support Software

Writing an interrupt handler in Turbo C is fairly straightforward because of the extensions to the language that permit control over 8086 device interrupts. The Turbo C extensions used by Commkit are described in the following paragraphs.

Table 2.4 Interrupt process identifiers.

Identifier	Interrupt Handler
`KEYIH`	Keyboard
`CLKIH`	Clock
`SP1IH`	Serial port 1
`SP2IH`	Serial port 2
`PPIH`	Parallel port

The `void interrupt` type can be associated with either a variable or a C function. A variable of this type can hold the 32-bit address (segment and offset) of an interrupt handler. A `voidinterrupt` function causes the compiler to generate the necessary instructions to save all registers on the stack upon entry to the function (it is assumed that the function will be called when an interrupt occurs, requiring the registers to be saved). Conversely, upon exit from the function, the compiler generates the code to restore the registers (by popping them from the stack). Instead of ending the function with a `ret` (return from subroutine), the compiler generates an `iret` (return from interrupt).

As an example, the keyboard interrupt handler entry point could be declared as follows:

```
void interrupt kb_ih()
{
/* Statements to handle keyboard interrupt */
}
```

`Setvect()` stores the address of an interrupt handler (which must be declared to be of type `void interrupt`) in the specified interrupt vector. For example, the code needed to store the address of `kb_ih()` (the keyboard interrupt handler) in interrupt vector location 9 (the keyboard interrupt vector, defined in `devices.h as KEY_INT`) is

```
setvect(KEY_INT, kb_ih);
```

The `getvect()` function returns a copy of the 32-bit interrupt address stored in a specific interrupt vector. For example, the code to obtain a copy of the original clock handler's address stored in interrupt vector 8 (defined in `devices.h as CLK_INT`) and to store it in the variable `old_clock` (declared to be of type `void interrupt`) would be

```
void interrupt (*old_clock)();
. . .
```

```
old_clock = getvect(CLK_INT);
```

The variable `old_clock` could be restored using `setvect()`:

```
setvect(CLK_INT, old_clock);
```

Interrupt Handler-Process Communication

Since different applications of Commkit may require that messages generated by the interrupt handlers be sent to different processes, the application software is expected to determine the destination of any interrupt-handler data. Accordingly, all interrupt handlers call the external function `low_level()` upon completion of their interrupt to allow the application to decide on the final destination of, for example, a character read from the keyboard. The major benefit of this approach is that it minimizes the need to modify `commkit.c`, thereby reducing the amount of recompiling necessary as well as decreasing the likelihood of introducing unexpected changes to `commkit.c`.

All Commkit applications must support an entry point to a function `low_level()`.

```
void low_level(int device, int code,
               unsigned char data)
```

where:

`device is` the process identifier of the device associated with the call to `low_level()` (one of `KEYIH, CLKIH, SP1IH, SP2IH`, or `PPIH`);
`code` is the cause of the call to `low_level()`, one of
`0`: The field is to be ignored.
`RECVDONE`: A data available interrupt has occurred (serial and parallel ports only).
`XMITDONE`: The transmit holding register is empty (serial ports only).
`MODEMSTATUS`: A modem status change has occurred (serial ports only).
`MSG_AVAIL`: A message has been received from a process for the interrupt handler specified in `device`. The message can be retrieved using `recv()`.
`data` is the data, if any, associated with the call to `low_level()`.

The Keyboard Interrupt Handler

The keyboard interrupt handler is an implementation of the algorithm discussed in Section 2.2.1, converting keyboard scan codes into their equivalent ASCII characters.

Since keys can be pressed in various combinations, the global state variables maintained by the keyboard interrupt are listed in Table 2.5. These

Table 2.5 Keyboard global-state variables.

Variable	Purpose
`left_shift`	State of LEFT SHIFT key
`right_shift`	State of RIGHT SHIFT key
`caps_lock`	State of Caps Lock key
`alt`	State of ALTMODE key
`CTRL`	State of CTRL key
`num_lock`	State of Num Lock key

state variables are all initialized to `FALSE`, indicating that the key has not been pressed. Once the key is pressed, the value changes to `TRUE`, returning to `FALSE` when the key is released. By maintaining this state information, it is possible to determine the various keys being pressed by the user; for example, CTRL-ALT-DEL can be detected if the variables `CTRL` and alt are `TRUE` and the character just assembled from the keyboard is DEL.

The keyboard interrupt handler software consists of three procedures:

`keyboard_init()`. The keyboard initialization software, responsible for initializing all keyboard state variables to `FALSE`.

`kb_ih()`. The keyboard interrupt entry point, called whenever a keyboard interrupt (i.e., make or break) occurs. The address of `kb_ih()` is stored in interrupt vector `0x09` (memory locations `0x24` through `0x27`). `Kb_ih()` passes control to `keyboard_handler()`.

`keyboardhandler()`. The keyboard-interrupt processing software, called from `kb_ih()` whenever a keyboard interrupt occurs. The keyboard handler uses the aforementioned state variables and two mapping tables (`shifted` and `unshifted`) to convert the scan code into the equivalent ASCII character and store it in the variable `ascii`. Once the character has been generated, the `keyboard_handler()` calls `low_level()` with the character `ascii` for forwarding to the appropriate process. If the combination of keys signal CTRL-ALT-DEL, execution is halted by the `keyboard_handler()` calling `rtn_to_dos`.

The Clock Interrupt Handler

The Commkit clock interrupt handler uses the 8253 timer chip (described in Section 2.2.1) to supply timing signals to applications. The clock initialization software does not change the clock counter value already defined by MS-DOS, for the following reasons:

In many versions of MS-DOS, the system clock is updated each time a clock interrupt occurs.

The floppy disk software uses the clock interrupt to determine when to switch off the disk motor; if the clock interrupt is blocked (i.e., the original clock interrupt handler is not called), once the motor starts it cannot be stopped (fun, but hard on the disk).

When a clock interrupt occurs, control is first passed to `clk_ih()` and then to `clock_handler()`. `Clock_handler()` then calls `low_level()` with an indication that a clock interrupt has occurred. The clock "ticks" `HZ` times a second, allowing `low_level()` to forward timing signals to various processes every 1/`HZ` seconds or perhaps to count interrupts and then to send a message after `HZ` interrupts have occurred (i.e., once a second).

Once the clock interrupt has been processed, control is *not* passed back to the previously executing process. Instead, the stack is modified to appear as if an interrupt has just occurred (the flags are pushed onto the stack), and the *original* MS-DOS clock interrupt handler is then called. This "faked" interrupt is achieved by calling `old_clock()`.

Turbo C "fakes" an interrupt by pushing the flags register onto the stack and calling `old_clock()` (which contains the address of the MS-DOS interrupt handler). To the called software, this looks like an interrupt stack frame. The return address is the location in `clock_handler()` after the call. This mechanism, whereby the Commkit application services the software associated with the interrupt and then calls MS-DOS, is known as chaining.

Once the MS-DOS clock interrupt handler has finished, the registers are restored to their original values and control is returned to `clk_ih()`. Since `clk_ih()` is declared to be of type `interrupt`, the stack contains the register values associated with the process that was running when the clock interrupt occurred: These registers are restored and control is returned to the original process.

Initializing and Restoring Interrupt Vectors

When any Commkit module begins execution, the interrupt vectors are, not surprisingly, associated with MS-DOS interrupt handlers. While the module executes, the keyboard, clock, and serial port interrupt handlers must refer to the Commkit interrupt handlers. When Commkit terminates, the interrupt vectors must be restored to the original MS-DOS interrupt handlers.

Procedure `initialize()` (found in `commkit.c`) is responsible for saving the MS-DOS interrupt vectors in five global variables: `old_clock`, `old_kb`, `old_spl`, `old_sp2`, and `old_pp`. These vectors, all of type `interrupt`, are obtained via `getvect()`. Upon termination, procedure `rtn_to_dos()` (in `commkit.c`) restores the five interrupt vectors using `setvect()`.

2.3.5
Message Handling

Commkit permits any entity (a foreground process, a background process, or an interrupt handler) to communicate with any other entity by sending messages. All processes are message-driven; that is, a process is idle until it is sent a message, at which point it is activated. Entities send messages using the `send()` primitive and receive messages using the `recv()` primitive.

Since an entity may be sent a message while it is processing an earlier message, all pending messages are queued until they can be received by the entity. Each Commkit entity is associated with its own message queue, as shown in Figure 2.8.

The message-handling software can determine the intended destination of a message because each message is sent with the process identifier of the destination entity. There is a one-to-one relationship between the entity, its queue, and its identifier number.

A message is any data structure having a maximum size of 128 bytes. When an entity sends a message, Commkit copies the contents of the data structure into the first available buffer taken from the queue `qe` (the contents of the message are ignored). The buffer is then kept on the queue associated with the destination entity. When an entity requests a message, Commkit checks the entity's queue. If the queue is not empty, the contents of the first buffer on the queue are copied into the message structure supplied by the destination entity. (An entity can determine the size of a data structure using the `sizeof` compiler directive.) Once read, the buffer is then returned to `qe`, the list of available buffers.

Commkit is responsible for supplying the message to the correct destination, so information must be maintained in addition to the message itself:

Information specific to the message (i.e., its size, `size`, and the identifier of the sending entity, `source`) is kept in the structure `queue_entry` (defined in `commkit.h`). The structure of a `queue_entry` is shown in Figure 2.9. Since there can be many messages waiting to be read by the entity, all queue entries associated with an entity are linked together through the field `next`.

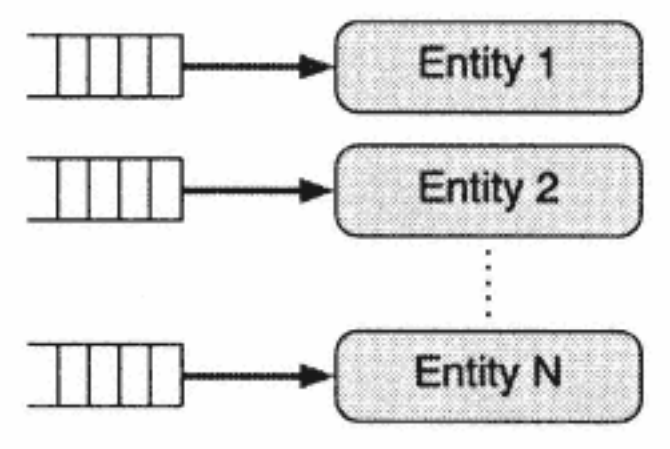

Figure 2.8
Commkit entities and their
message queues.

Figure 2.9
The structure of a queue entry.

Since there are many entities and each has its own list of messages, Commkit maintains a table `queue`, in which each entry corresponds to the list of messages waiting to be received by an entity (the identifier is used as a subscript to access an entity's queue). An entry in the table `queue` consists of two pointers: `head`, which points to the entity's first pending message, and `tail`, which points to the last pending message. In addition, each queue entry has a field `count`, which indicates the number of pending messages waiting to be received by the entity.

In Figure 2.10, three messages are pending for entity number `7`: The first is a single byte *S* sent from entity `3`; the second is a 10-byte message sent from entity `9`, and the third is a 2-byte integer (with value `0`) from entity `1` (`head` points to the first available message, while `tail` points to the last message in the list).

A second set of data structures is maintained to handle messages intended for background processes. `Num_back` is a global count used by Commkit to indicate the number of background processes. Initially `num_back` is 0; however, it is incremented as each background process is registered, using `register_background()`. In addition, the array `back_list` contains the list of background process identifiers and entry points for each background process. This list is used in `do_other_tasks()` to determine which background process has a pending message and, if a message is available, the address of the procedure (i.e., the process entry point).

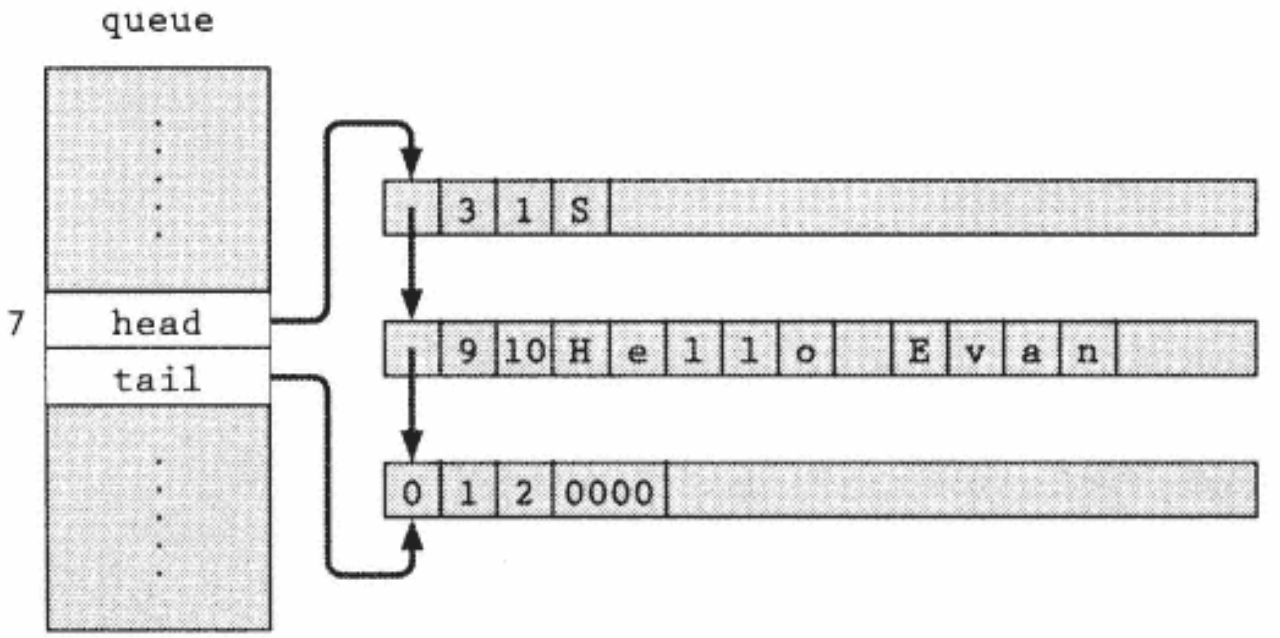

Figure 2.10
A message queue for a Commkit entity.

The Send Primitive

The send primitive allows an entity to send a message to another entity. The transmitting entity must specify the intended destination's identifier, the message, and the message size. The send software stores the message on the queue associated with the supplied destination identifier.

The format of the send primitive is as follows:

```
int send(int src, int dst, char *msg, int size)
```

where:

`src` is the identifier of the transmitting (i.e., source) entity.
`dst` is the identifier of the intended destination entity.
`*msg` is the address of the message to be transmitted to the destination entity. The address is specified by prefixing the data structure with &. The message is stored on the queue associated with the destination entity (as indicated by the `dst` identifier).
`size` is the size (in bytes) of the message `msg`.

For example, the procedure for the foreground process `APPLICATION` to send a 64-byte message, `msg`, to the background process `BACKGROUND_1` could be written as follows:

```
send(APPLICATION, BACKGROUND_1, msg, 64);
```

The Receive Primitive

The receive primitive allows an entity to receive a message sent from another entity. The receiving entity specifies the maximum number of bytes that it is willing to accept and supplies the address of the buffer into which the message should be copied. The number of bytes the receive primitive copies into the message buffer is the *smaller* of either the size of the message (as specified by the source) or the number of bytes that can be accepted by the destination.

The format of the receive primitive is

```
int recv(int *src, int dst, char *msg, int size)
```

where:

`*src` is the identifier of the entity that has sent the message; returned by the receive primitive software. The identifier can be returned, since the address of `src` is supplied (using &).
`dst` is the identifier of the receiving entity. The receive primitive software uses this identifier to determine the queue from which to remove the message.
`*msg` is the address of a data structure where the message is to be stored.

The address is specified using &.

`size` is the size (in bytes) of the message. The entity must specify the *maximum* size the data structure can accept. The receive software compares the two sizes associated with the message and chooses the *smaller*.

The number of bytes copied into the message buffer is returned by the `recv()` function.

As an example, consider the following code fragment, which allows the foreground process, `APPLICATION`, to receive a message from any other source entity:

```
void receive_from_any()
{
int src;
char data;
if (recv(&src, APPLICATION, &data, 1) == 1)
    /* Code to handle a one-byte message */
}
```

In some situations, a receiving entity might be required to wait for messages from a number of different transmitting entities. If the messages are of different sizes, the receiving entity must specify the *largest* size possible to ensure that no data is lost when it receives a message. For example, in the following code fragment, the process `BACKGROUND_1` can accept messages 9 bytes in length from process `BACKGROUND_2` and 1-byte messages from process `APPLICATION`:

```
void special_receive()
{
int src;
char msg[9];
recv(&src, BACKGROUND_1, &msg, 9);
switch (src)
{
case BACKGROUND_2:
   /* Data in msg[0] through msg[8] */
   break;
case APPLICATION:
   /* Data in msg[0] */
   break;
default:
   message(0, 0, INVERSE,
           "Bad message in BACKGROUND_1");
}
}
```

2.3.6
Implementation of the Message-Handling Software

Interprocess communication via messages in Commkit involves both the transmitting and receiving entities, as well as the procedures `send()` and `recv()`.

Send() Implementation

The `send()` function is essentially an enqueuing operation; the message supplied by one entity is placed on the queue of the destination entity. Since Commkit permits both "high-level" software and interrupt handlers to send messages, it is possible for an interrupt to occur at the same time a highlevel process is sending a message (i.e., enqueuing the message). Should the interrupt handler also send a message, the queue data structures may be damaged, resulting in inconsistencies. Sections of code that access shared data (such as the message queue) are known as *critical regions* and must be protected against "concurrent" access.

In Commkit, the queue is a critical region, and it is protected from concurrent access by the disabling of interrupts. The `send()` function does not disable interrupts explicitly; instead, all the parameters (i.e., `src, dst, *msg`, and `size`) are placed in a structure `msg_struct` (defined in `commkit.h`). The address of `msg_struct` is passed to Commkit using the `ds` register and causing a software interrupt via interrupt vector 50. Software interrupts are identical to hardware interrupts in that the flags register, the code segment register, and the instruction pointer are all pushed onto the stack; interrupts are also disabled. (The `send()` function is implemented in `srcalls.c`.)

An interrupt handler, `send_ih()`, is associated with interrupt vector 50 and implemented in `commkit.c`. The `ds` register (still pointing to `msg_struct`) is used as a pointer to the various fields in `msg_struct`; these fields are extracted and a call takes place to `low_level_send()`, which is responsible for enqueuing the supplied message. Once enqueued, control returns to `send_ih()`, which, in turn, returns control to `send()`. The `send()` cycle is shown in Figure 2.11.

Once a message has been enqueued, it can be received.

Recv() Implementation

When an application attempts to receive a message, the flow of control is, for the most part, the same as that found in a message transmission. That is, the application calls `recv()`, which causes a software interrupt to `recv_ih()` with the parameters (`*src`, `dst`, `*msg`, and `size`), all placed inside a structure of type `msg_struct`. The `ds` register points to this structure. `Recv_ih()` then calls `low_level_recv()`, which attempts to dequeue the first available message for the application. At this point, the actions that take place depend upon the destination of the message (specified in the field `dst`):

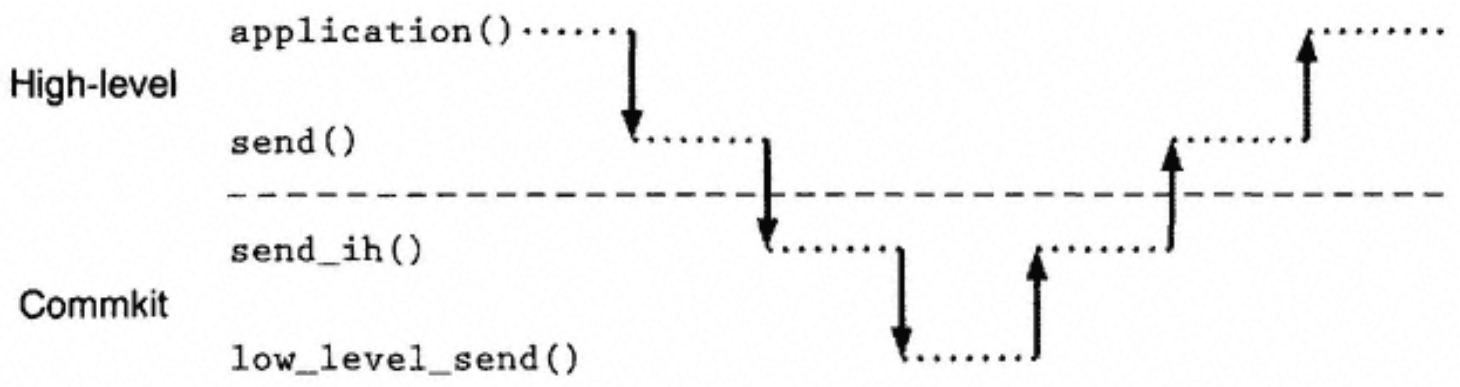

Figure 2.11
The `send()` cycle.

Foreground. If a message is available on the foreground process's queue, when the foreground process executes a `recv()`, the message is returned immediately. However, if the foreground process's queue is empty, the foreground process is suspended and control passes from `recv()` to `do_other_tasks()`. `Do_other_tasks()` performs other tasks (see below) until a message becomes available for the foreground process. When a message is placed on the foreground process's queue, control returns to `recv()`, which copies the data from the queue into the foreground process's message, and execution of the foreground process resumes.

Background. The background process is activated when the foreground process is suspended and a message is available for the background process. The background process is called from `do_other_tasks()`. `Do_other_tasks()` obtains the process identifier and the entry point from the array `back_list`. The background process continues to be called as long as messages exist for it and the foreground process is suspended. The background process should have a call to `recv()` so that the pending message is read immediately. If the background process fails to perform a `recv()`, it will continue to be called since its message queue is never emptied!

Interrupt handler. A message arriving for an interrupt handler results in a call to `low_level()` from `send()` with a `code` of `MSG_AVAIL`. If the interrupt handler can accept the message, it is expected to perform a `recv()`; otherwise, the call is ignored (presumably the message is received at a later time).

The `recv()` function, like the `send()` function, is implemented in `srcalls.c`.

Miscellaneous

The `send()` and `recv()` primitives call the interrupt handlers `send_ih()` and `recv_ih()` via interrupts 50 and 51, respectively. The interrupt vectors are initialized in procedure `initialize()` by first saving the original interrupt vectors in `old_vec50` and `old_vec51` (using `getvect()`) and

then assigning the vectors the values of `send_ih()` and `recv_ih()` (using `setvect()`). As with all interrupt vectors used by Commkit, they are restored by a call to `rtn_to_dos()`.

When a call is made to `low_level()` and a message is to be sent or retrieved, it is not necessary to call `send()` or `recv()` since interrupts are already disabled. Instead, procedures `low_level_send()` and `low_level_recv()` can be called directly. This reduces the repeated stacking of the registers whenever an interrupt occurs.

For example, to send a 1-byte message from the clock interrupt handler (`via low_level()`) to a background process `BACKGROUND_3`, one could write:

```
void low_level(int device, int action, char data)
{
if (device == CLKIH)
    send(CLKIH, BACKGROUND_3, &data, 1);
  }
```

2.3.7
Support Routines

In addition to the transmission and reception of messages, there are a number of functions offered by Commkit that can be used for support purposes such as screen output, numeric conversion, and screen clearing. These routines are intended for use by the interrupt handlers, although they can be employed by either the foreground or the background processes. They are found in `commkit.c` and described in the following paragraphs.

To display a single `character` on the screen at a specific `line` and `column`, the `display()` procedure is used:

```
void display(int line, int column, char character,
             char attributes)
```

The `attributes` indicate how `character` is to be displayed. Attribute values defined in `devices.h` include `NORMAL` (white on black), `INVERSE` (black on white), and `BLINK` (flashing on and off). Note that the `BLINK` attribute must be or'ed with another attribute in order to function. For example, the following code displays the uppercase letters of the alphabet across the middle of the screen (line 12, columns 27 through 52) with attributes `BLINK` and `INVERSE` set for each letter:

```
void display_example()
{
int i;
char ch;
```

```
    i = 27;
    for(ch = 'A'; ch <= 'Z'; ch++)
        display(12, i++, ch, INVERSE I BLINK);
    }
```

`Message()` displays a `string` with certain `attributes` starting at a specific `line` and `column`:

```
    void message(int line, int column, int attributes,
                 char *string)
```

`String` must be null-terminated ('`\0`'). The same attributes apply to `message()` as to `display()`. The following example shows how a string can be displayed in the upper left-hand corner of the screen:

```
    message(0, 0, NORMAL, "Press ENTER to continue");
```

The above string is null-terminated since, by default, C appends the null character to the end of any text string (i.e., a string enclosed in double quotes).

A 16-bit quantity (such as an integer) can be converted to a 5-byte string using `int2hex()`:

```
    void int2hex(int number, char *string)
```

String must be at least five bytes long to ensure that sufficient storage exists for four hex digits *and* a null character, `'\0'.Int2hex()` is generally used in conjunction with `message()`. The following example shows how the contents of the variable `data` can be displayed:

```
    void hexoutput(int data)
    {
    char mask[5];
    int2hex(data, mask);
    message(10, 0, INVERSE, "Data: ");
    message(10, 7, INVERSE, mask);
    }
```

The cursor can be moved to a specific location on the screen using `move_cursor()`:

```
    void move_cursor(int newx, int newy)
```

For example, to move the cursor to the center of the screen (x: 12, y: 40):

```
    move_cursor(12, 40);
```

This routine does not work on all PCs because of differences between various video cards.

The message queues can be dumped at any time by calling

```
void dump_queues(int line)
```

Each queue entry is a pointer to a message (maximum length 128 bytes). `Dump_queues()` displays the queue entries associated with each entity (i.e., process or interrupt handler) as a pair of bytes, starting on the specified `line` of the screen. The first byte displayed (in inverse video) is the identifier of the entity to which the messages are destined. Subsequent bytes are displayed in pairs: the first indicates the source of the message (in normal video), and the second contains the first byte of the message (in inverse video).

To save space, empty queues are not displayed.

Portions of the screen can be cleared using

```
void clear_scr(int xup, int xlo, int yup, int ylo)
```

The screen is cleared between points `xup`, `yup` and `xlo`, `ylo`. The following code fragment shows how the lower half of the screen (lines 12 through 25) can be cleared:

```
clear_scr(12, 25, 0, 79);
```

2.3.8
Co-existing with MS-DOS

Unlike many applications that run in an MS-DOS environment, Commkit controls a number of the PC's interrupts. Control over interrupts is necessary in order to understand how hardware and software interact to support data communications. There are two exceptions to this: First, MS-DOS and BIOS disk access software are still permitted to function (since there is no single standard of disk supplied with the PC), and second, clock interrupts are chained from Commkit to MS-DOS.

When an interrupt occurs that passes control to MS-DOS, MS-DOS changes two sets of registersthe 32-bit stack pointer (consisting of a 16-bit segment and a 16-bit offset register pair) and the 16-bit data segment registerto point to those segments specific to MS-DOS. Most, if not all, MS-DOS interrupt handlers enable interrupts soon after entry into the interrupt handler. Upon completion of the interrupt, the original stack pointer and data segment registers are restored and control returns to the interrupted procedure.

Although control will pass to the correct Commkit interrupt handler should an interrupt occur from a device controlled by Commkit, when an MS-DOS interrupt handler is active (remember, the code segment-offset is

taken from the interrupt vector) there are two important points that must be taken into consideration:

The data segment register will be pointing to the data segment associated with MS-DOS, not Commkit.

The MS-DOS stack *may* not be large enough to accommodate the stack requirements of the Commkit interrupt handlerpossibly resulting in stack overflow and the destruction of part of the MS-DOS data area.

Fortunately, part of the problem is already solved: when an interrupt occurs, the data segment register is pushed onto the stack (thereby saving the MS-DOS data segment register value) and Turbo C assigns Commkit's data segment to the data segment register. When the interrupt handler has finished, the data segment register is popped from the stack, thus restoring the register to MS-DOS's data segment.

To overcome the potential stack problems, the following algorithm is implemented in each interrupt handler:

1. Upon entry to the interrupt handler, save the old stack segment and stack offset registers.

2. Set the stack segment and stack offset to the local stack associated with the interrupt handler.

3. Perform the task required to service the interrupt.

4. Restore the old stack segment and stack offset to permit control to return to the interrupted software.

Changing the stack pointer requires access to the 8086 processor's stack segment register and the stack offset register, both of which are available to Turbo C as the pseudo-registers `_SS` (stack segment) and `_SP` (stack offset). (The pseudo-registers are mapped into the actual registers by the compiler; they are declared `unsigned int`.) Saving and restoring the stack pointer is achieved by

```
void interrupt sample_ih()
{
 /* Saving. . . */
old_ss = _SS;
old_sp = _SP;
. . .
/* Restoring. . . */
_SS = old_ss;
_SP = old_sp;
}
```

The variables `old_ss` and `old_sp` cannot be declared as variables local to the interrupt handler, since they will be stored on the stack set up upon entry (i.e., the original stack). If the stack pointer registers (`_SS` and `_SP`)

are subsequently changed, `old_ss` and `old_sp` will no longer refer to their locations on the original stack, meaning that the original stack pointer will be lost. Therefore, `old_ss` and `old_sp` must be declared as globals.

Fortunately, changing the stack pointer registers is a straightforward operationit simply involves setting the register to an array stored globally (i.e., in the interrupt handler's data segment):

```
char lcl_stk[256];

void interrupt sample_ih()
{
/* Old _SS and _SP saved */

_SS = _DS;
_SP = (unsigned int) &lcl_stk[255];

/* Old _SS and _SP restored */
}
```

The stack segment pseudo-register, `_SS`, is assigned the value of the data segment pseudo-register, `_DS`, ensuring that the segment is that of `lcl_stk`. The stack offset pseudo-register, `_SP`, is then assigned the address of the last location in the stack (since stacks operate from high to low memory).

Each Commkit interrupt handler has its own stack, declared of type `stack` (defined in `commkit.h` and used in `commkit.c`):

```
struct stack
{
unsigned int ss;
unsigned int sp;
char stk[256];
};
```

2.3.9
Miscellaneous

All modules require the following header files to permit compilation to proceed:

```
#include "general.h"
#include "ascii.h"
#include "devices.h"
```

The header file `general.h` contains a list of common directives; `ascii.h` is the ASCII character set; and `devices.h` contains the definitions of various device ports. `General.h` also includes the definitions for `send()`, `recv()`, `low_level()`, and `register_background()`.

2.4
Interprocess Communication (IPC) Using Commkit

To illustrate how Commkit supports interprocess communication (or IPC), consider the following problem:

Write a program that will allow a user to enter data from a PC's keyboard and display it on the PC's screen. Simultaneously, the time elapsed since the program started execution should be displayed as `'hh: mm: s s'` in the upper left corner of the screen. When the user attempts to type beyond the last line of the screen, the screen should be cleared and the cycle should continue. A CTRL-C entered by the user should terminate the program, returning to MS-DOS.

2.4.1
Design Considerations

`Commkit.c` is supplied with software that handles clock interrupts, keyboard interrupts, and screen character output. However, software must be written that meets the problem's other criteria, notably clearing the screen, watching for CTRL-C, and formatting the clock information into hours, minutes, and seconds.

Instead of writing a single process that handles all of these functions, the problem is divided into three separate tasks to illustrate how Commkit supports interprocess communication:

1. A process to examine each character entered by the user to determine whether it is a CTRL-C.

2. A process to accept a 1-second timing signal from the clock and update the current time.

3. A process to display the clock information and the keyboard input on the screen.

Figure 2.12 illustrates the different paths that the information can take, where:

P1 is a process that accepts characters from the `keyboard_handler()`, scans for CTRL-C, and forwards all other characters to P3.
P2 is a process that accepts timing signals from the `clock_handler()`, and converts them into a data structure consisting of hours, minutes, and seconds. The data structure is then forwarded to P3.

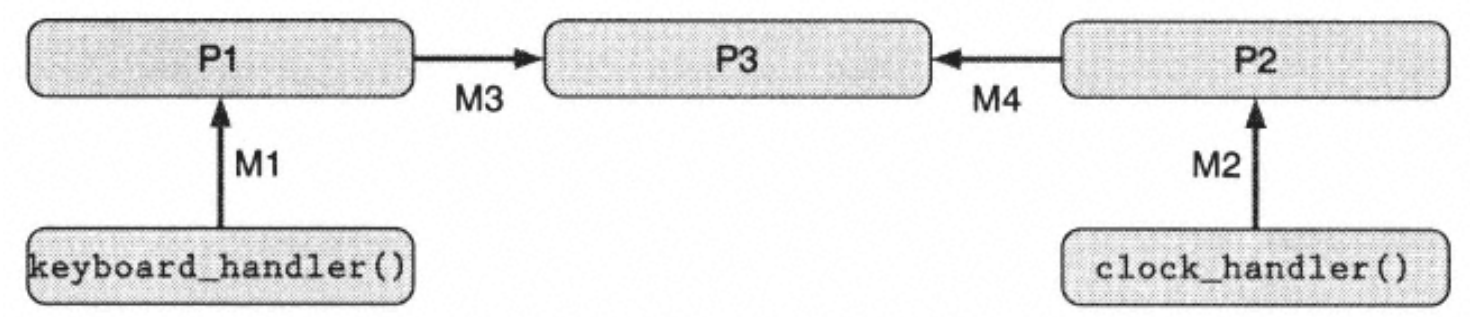

Figure 2.12
Information flow for the IPC example.

P3 is a process that displays individual characters (from P1) or character strings (from P2). This process is also responsible for clearing the screen should printing an individual character cause it to go off screen.
M1 is a character sent in a message from the `keyboard_handler()` to P1.
M2 is a 1-second timing signal message sent from the `clock_handler()` to P2.
M3 is a character (other than TRL-C) sent in a message from P1 to P3.
M4 is the current time string (expressed as `hh:mm: ss`) sent from P2 to P3.

2.4.2
Implementation

An example of an implementation of the problem is available as the file `ipc.c` on the Commkit distribution diskette. The file consists of three processes (one foreground and two background), described in the following paragraphs.

`do_ipc(). Do_ipc()` is the foreground process (identifier `APPLICATION`); it accepts single-byte messages sent from the keyboard handler (identifier `KEYIH`). When a CTRL-C (*ETX*) is detected, `do_ipc()` terminates, returning to the mainline (and eventually MS-DOS). All other characters are sent to process `display_screen()` (identifier `BACKGROUND_2`).

`clock_signal(). Clock_signal()` is a background process (identifier `BACKGROUND_1`) that receives timing signal messages once every second from the clock interrupt handler, `clock_handler()`. Upon receipt of the message, the number of seconds is increased, the number of minutes is increased (if seconds exceed 59), and the number of hours is increased (if minutes exceed 59). The current hours, minutes, and seconds are then stored as a string of characters in the structure `tod` and sent to process `display_screen()` for display.

`display_screen(). Display_screen()` (identifier `BACKGROUND_2`) accepts either single-or multiple-byte messages. Single-byte messages (from `do_ipc()`) are displayed in the next screen position, indicated by the values of `line` and `column`; multiple-byte messages are assumed to be clock messages (from `clock_signal()`) and are displayed in the upper left corner of the screen.

Two static variables are maintained by `display_screen()`. Column indicates the present horizontal character position, and line contains the current line number. When `column` exceeds 79 or a carriage return is detected, `column` is set to 0 and the line number is incremented. When `line` exceeds the screen length, the screen is cleared and the next character is displayed at the top of the screen.

Before the program can be run, the background processes must be registered and the low-level software must be written:

The background processes are registered with Commkit in the mainline of `ipc.c`:

```
register_background(BACKGROUND_1, clock_signal);
register_background(BACKGROUND_2, display_screen);
```

The Commkit interrupt handlers use `low_level()` to communicate with the foreground and background processes. In `ipc.c, low_level()` is called from either the clock handler (`clock_handler()`) or the keyboard handler (`keyboard_handler()`). Characters from the keyboard are sent immediately as messages to the foreground process. One-second timing signals are sent to the background process associated with identifier `BACKGROUND_1` after `HZ` interrupts have occurred; the static variable `sec_count` keeps track of the number of clock interrupts that make up one second.

2.4.3
A Software Walk-through

Internally, the steps taken to display a character on the screen once it has been entered are as follows:

1. The `keyboard_handler()` receives a series of interrupts indicating that a key has been pressed and released. Once the value of the character has been determined and stored in the variable `ascii`, the `keyboard_handler()` calls `low_level()`, which sends the character to `APPLICATION`:

```
low_level_send(KEYIH, APPLICATION, &data, 1);
```

2. The keyboard character is kept on the queue `APPLICATION` until the foreground process performs a `recv()` with identifier `APPLICATION`.

3. When the foreground process `do_ipc()` performs a `recv()`, the next available character from queue `APPLICATION` is returned:

```
recv(&src, APPLICATION, &ch, 1)
```

The byte received, `ch`, is then checked. If `ch` is an *ETX*. control is returned to MS-DOS; otherwise, the byte is sent to process `BACKGROUND_2`:

```
send(APPLICATION, BACKGROUND_2, &ch, 1);
```

`Do_ipc()` then executes another `recv()`, waiting for its next message.

Clock interrupts follow a similar path from the `clock_handler()`, via `low_level()`, to `clock_signal()`. However, unlike the keyboard interrupts, clock interrupts occur on a regular (1-second) basis, signaling that the current hours, minutes, and seconds are to be converted into a string of the

format `hh:mm:ss`. The string is then sent to `display_screen()`, which is responsible for displaying it. Note that since `clock_signal()` is a background process (identifier `BACKGROUND_1`), it is executed only when the foreground process is suspended.

Process `display_screen()` is activated whenever a message is detected on queue `BACKGROUND_2` and the foreground process is suspended. When called, `display_screen()` reads the message into a 10-byte buffer using `recv()`:

```
recv(&src, BACKGROUND_2, &msg, 10);
```

There are two possible sizes of message that `displayscreen()` can receive: a single-byte character sent from `do_ipc()` and a time string sent from `clock_signal()`. Although `clock_signal()` can distinguish between the messages by the message's source identifier, it uses instead the message size: Single characters are displayed directly in the next available screen position; strings are displayed in the upper left corner of the screen.

The message flow between the interrupt handlers and the processes is illustrated in Figure 2.13.

2.5 Compiling Commkit

The Commkit software is written in Turbo C, using the Turbo C extensions for handling interrupts. As described in Appendix A, the `make` utility ensures that only those modules that have been changed are actually recompiled and linked. All modules can be compiled and linked without the use of Turbo Assembler.

Assuming that changes have been made to `ipc.c`, the new `ipc.exe` executable is created by typing:

```
C:\> make ipc.exe
```

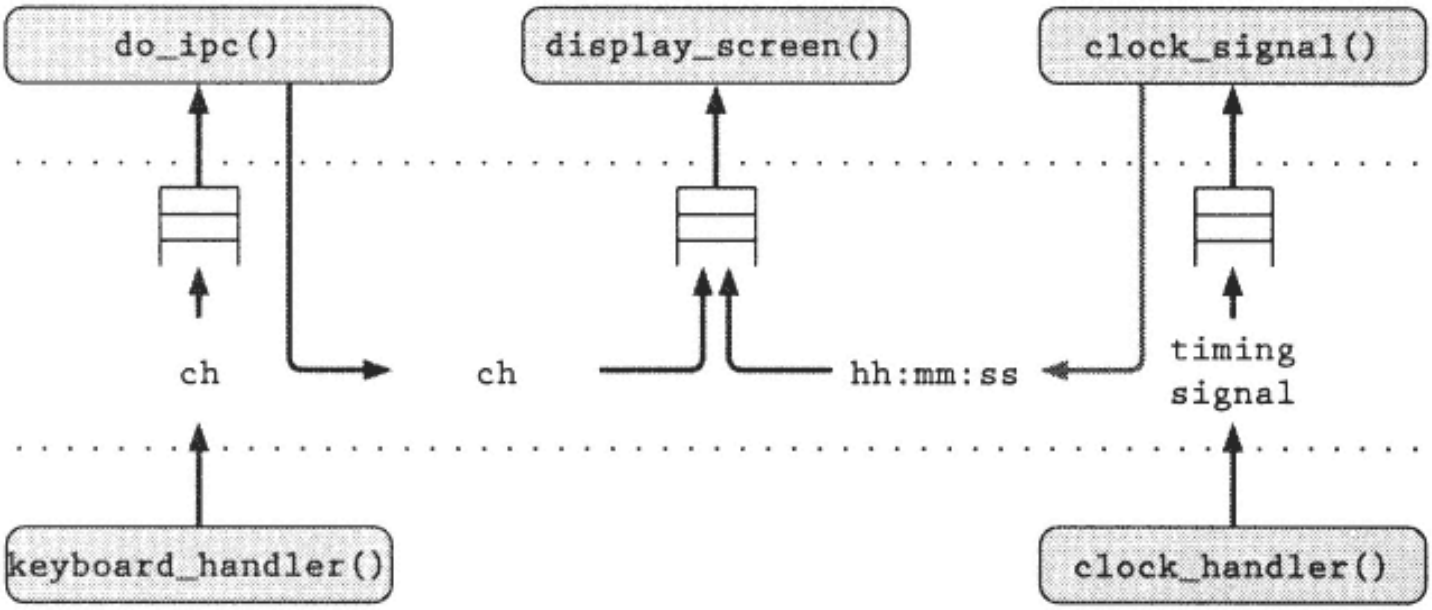

Figure 2.13
Message flow in the IPC example.

`Make` will detect that `ipc.c` is "younger" than `ipc. obj` and will proceed to compile `ipc.c` into a new object module using the options `-ms` (small model) and `-c` (create an object file only). All of the necessary include files are accessed by the compiler. The Turbo linker, `tlink`, is then invoked because make detects that `ipc.obj` is "younger" than `ipc.exe`. The linker links `ipc.obj` and `commkit.obj` as well as the small-model libraries to form `ipc.exe`.

To run the compiled program, simply type

```
C:\> ipc
```

At this point, the screen will be cleared and the elapsed time since the start of execution will appear in the upper left of the screen. Whenever a character is typed, it will appear on the screen. Lines can be terminated explicitly using carriage returns or implicitly when the line is filled (i.e., 80 characters have been typed on a line); in either case, output is started on the next line. If the screen is full, the next character entered will cause the screen to clear and the character will be placed in the first column on the line under the elapsed time.

2.5.1
When Things Go Wrong

On occasion, programs don't do what the programmer expects them to do, requiring the programmer to halt program execution. Any Commkit program can be aborted by typing CTRL-ALT-DEL; this is *not* the MS-DOS reboot, it is simply a signal to Commkit to restore the original interrupt vectors and return control to MS-DOS. All of the user-level processes supplied with Commkit can also be stopped using CTRL-C.

Software errors can be traced by placing `display()`, `message()`, or `printf()` procedure calls throughout the code. Since `printf()` is not recursive, only `display()` and `message()` should be used in interrupt handlers. The `dump_queues()` procedure can also be used, usually as a last resort.

A process can send a message to any other process, with two possible results:

1. If a message is sent to the foreground process, a background process, or an interrupt handler, the destination process will eventually be called.

2. If a message is sent to a nonexistent process, `send()` will fail, returning a value of-1.

Should the destination process or interrupt handler not read its queue whenever a message becomes available, Commkit can run out of message space. If this occurs, the message queues are dumped and control is returned to MS-DOS via `rtn_to_dos()`.

2.6 Summary

This chapter introduced the tools that will be used throughout the remainder of the book: the PC and Commkit.

The PC devices are all interrupt-driven and controlled by the Commkit software. Interrupts are converted into messages that are sent from the interrupt handlers (calling `low_level()`) to processes.

Commkit supports two types of process (in addition to the interrupt handlers): foreground and background. There is one foreground process, associated with identifier `APPLICATION`, but there can be 0 or more background processes, each with its own unique identifiers.

Messages are sent using the `send()` primitive and received using the `recv()` primitive. The intended destination of a message is indicated by the destination identifier specified by the transmitting process.

Finally, it is worth noting that communications within the DTE itself are subject to protocols. For example, certain rules must be followed in order for the DTE to receive and accept a character from the keyboard. Similarly, when using the interprocess communication primitives, rules must be followed if the data being sent is to be received and processed correctly.

2.7 Exercises

1. (B) Commkit maps the keyboard scan codes into an equivalent set of ASCII characters. Use `ipc` to experiment with the various key combinations (i.e., CTRL, ALTMODE, and SHIFT) to determine what characters are displayed when the various keys are pressed.

2. (B) Does the "digital" clock displayed in the upper left corner of the screen keep accurate time when ipc is running? Devise a method to determine the accuracy of the clock.

3. (B) Is the accuracy of the clock affected by the typing of characters when `ipc` is running? Explain.

4. (C) Explain what the following foreground process displays and whether the data displayed ever changes. What process supplies the initial input? It may be assumed that Commkit is functioning correctly and that the foreground process receives all keyboard input.

```
void do_fun()
{
int src;
char ch;

while (TRUE)
{
      recv(&src, APPLICATION, &ch, 1);
```

```
        display(12, 40, ch, NORMAL);
        send(src, APPLICATION, &ch, 1);
    }
    }
```

This question leads to some interesting problems; try implementing the process to see what actually happens.

5. (C) It is often useful to know what a program does when it crashes. Modify `ipc.c` so that the foreground process `do_ipc()` never reads the messages that are sent to it. What happens when `ipc` is run and characters are entered at the keyboard? Continue typing for about 30 secondswhat happens now? Why doesn't CTRL-C cause Commkit to return to MS-DOS? Does the clock signal appear in the upper left corner? Explain. If the suspense gets too great, CTRL-ALT-DEL will return control to MS-DOS.

6. (C) Add character deletion to `ipc.c` using backspace to indicate the deletion.

7. (C) Modify `ipc.c` so that the user can treat the screen as an edit buffer, using the up, down, left, and right arrows to move throughout the text.

8. (C) Modify `ipc.c` so that the user can enter the current time of day. When a special character (such as CTRL-A) is entered, this should signal `do_ipc()` that a new value for the time of day is about to be supplied. The next 8 characters (in the format `hh/mm/ss`) can be treated as the new time. Remember that `KEYIH` supplies one character at a time.

9. (C) Rewrite `dump_queues()` so that complete messages will be displayed.

10. (C) Add another background process to `ipc.c` that takes the output from `do_ipc()` and changes the case on each letter before forwarding it to `display_screen()`. The information flow should be as shown in Figure 2.14.

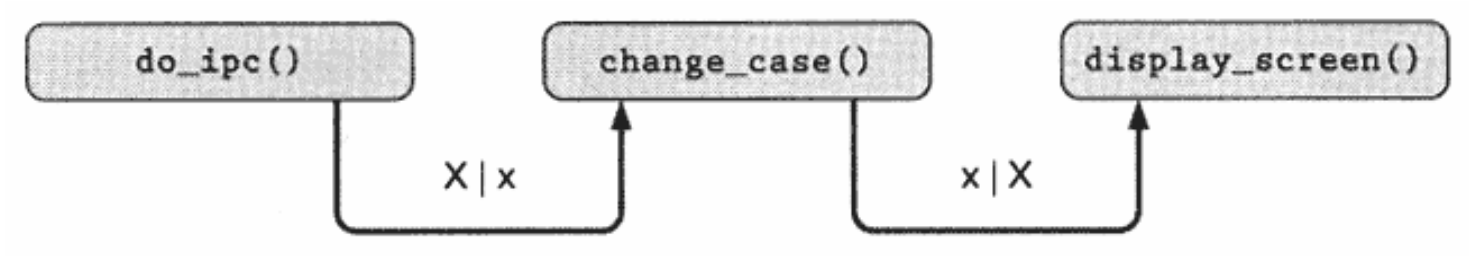

Figure 2.14
The `change_case()` process.

The new process (`change_case()` in Figure 2.14) should receive any message sent to queue `BACKGROUND_3`. Once the case has been changed, the character should be forwarded to queue `BACKGROUND_2`.

PART 2
POINT-TO-POINT COMMUNICATIONS

Most computer communications occur between pairs of DTEs. One widespread application of computer communications involves the transfer of information from an individual's terminal to a remote computer. This is commonly known as an asynchronous point-to-point communication.

In Chapter 3, asynchronous point-to-point communications are discussed in light of

The general concepts and terminology associated with asynchronous communications.

The hardware required to support asynchronous communications (illustrated using the PC's asynchronous hardware).

The Commkit software that controls the asynchronous hardware.

Communications between DTEs cannot occur without some form of connection. Chapter 4 examines some of the types of media and standards that are available to support communications between pairs (or groups) of DTEs.

Data communication software, like any other type of application software, may be subject to errors in design or implementation. Software debugging techniques, such as inserting `printf()` statements within the code to display values, often lead to further problems (the time taken to display values may result in the loss of data from the communications channel). Instead, what is needed are tools to permit the "eavesdropping" of the traffic on the channel. In Chapter 5, techniques for analyzing the contents of the communications channel are examined and a Commkit tool for channel monitoring is presented.

3
Asynchronous Communications

3.1
Introduction

An asynchronous communication is one that can occur without regard to time. Asynchronous communications are typified by terminalhost communications in which the user can send data from the terminal to the host periodically at random, unpredictable intervals. Due to the widespread popularity and availability of asynchronous communications (for example, the PC is equipped with at least one asynchronous communications adapter), this chapter examines them in detail.

3.2
General Concepts

In Chapter 1, information within a DTE was shown to be stored in binary form. For example, the ASCII letter S is represented by the 7-bit binary pattern:

```
1010011
```

When information is transmitted on the channel between DTEs, it is not sent as a letter (such as *S*), but rather in a form that can be easily manipulated by the DTEs. Since the information is already stored in binary within the DTE, data is, not surprisingly, transmitted as a series of `0`s and `1`s. As an example, the transmission of the letter *S* is represented pictorially in Figure 3.1.

When data is transferred between DTEs, it is usually done so serially, since it is cheaper to use one wire instead of the seven or eight (depending upon the number of bits in the byte) that would be required in a parallel transmission. The rightmost (low-order) bit is sent first, followed by the remaining bits, working from right to left.

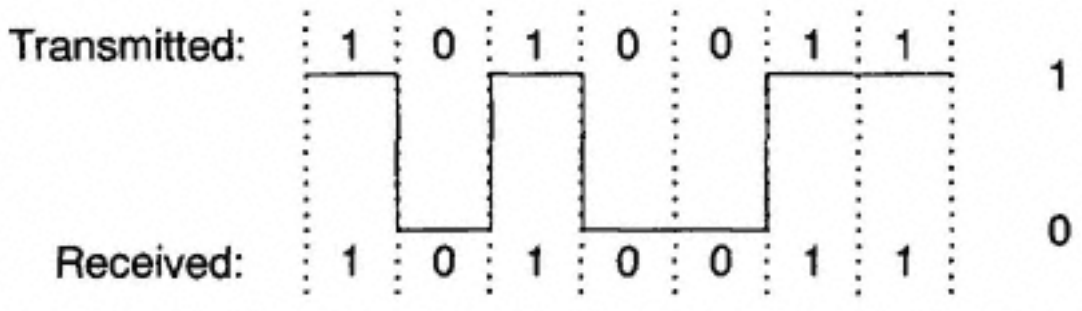

Figure 3.1
Transmission of the letter *S*.

It is generally accepted that a transmitted 1 is called a mark, while a transmitted `0` is called a space. These terms come from the days of telegraphy when a 1 on the communication channel was represented by a flow of current that resulted in the equipment making a mark on the receiver's output device (a moving strip of paper) while the lack of current left a space.

The receiving DTE must read the data bits from the communication channel at the same rate at which they are being transmitted; otherwise, a loss of information will occur. Figure 3.2 depicts what might result if the byte *S* is received at twice the rate of transmission; the receiver receives *two* bytes instead of one.

Similarly, if the transmitter is transmitting faster than the receiver can accept, there may also be a loss of information. In Figure 3.3, the transmitter is sending data at twice the rate it is being received (the letter *S* is sent once again). The receiver receives only four bits rather than the seven that make up the byte. The value of the third bit could be either a mark, 1, or a space, `0`, depending upon the exact moment the channel was sampled (channel sampling is discussed further in Section 3.3.2).

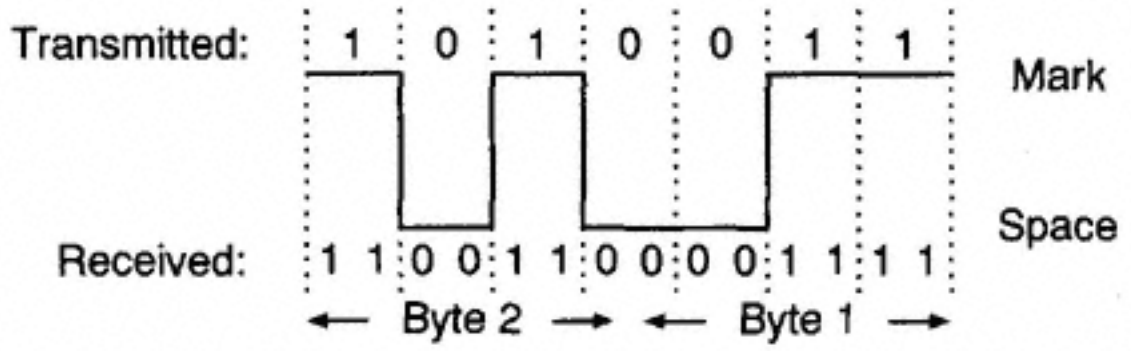

Figure 3.2
Receiving at twice the rate of transmission.

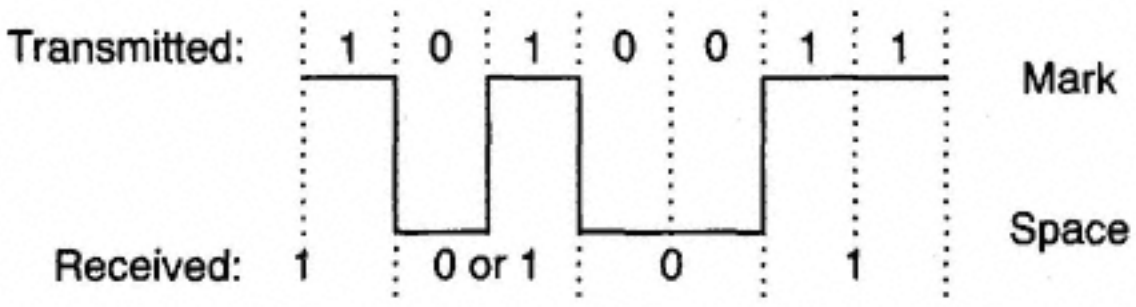

Figure 3.3
Receiving at half the rate of transmission.

As the previous examples illustrate, in any communication both the transmitting and receiving DTE must agree to a single, common channel speed, known as the line speed.

Asynchronous line speeds can be expressed in terms of bits per second (or bps) and typically range from 50 bps to 9600 bps. Some commonly used line speeds are 1200 bps, 2400 bps, and 9600 bps. As well as agreeing upon the line speed, both DTEs must concur upon the number of bits in a byte (this is usually a choice between 7-bit and 8-bit, as these are the most common character code widths).

Since the channel can be only in one of two states, mark or space, it must be in one or the other when there is nothing to transmit. All DTEs, when they are not transmitting, leave the channel in the mark state.

This then raises another issue: if the channel idles in the mark state, how does the receiving DTE determine where the first data bit occurs? In the worst case, if a DTE transmits the binary pattern `1111111', the receiving DTE will not be able to distinguish between the marks representing the idle channel and the marks representing the data!

The solution to this problem is to have the first bit sent as a "nondata" space. That is, before the first data bit is sent, the channel changes state from mark to space. After the space, the bits making up the byte are transmitted. Figure 3.4 shows the transmission of the byte *S*. The first bit, indicating the start of data, is called the start bit. The start bit signals the receiving DTE that a data byte is about to follow. The width of the start bit is the same as that of each of the data bits (that is, it takes the same length of time to send the start bit as it does to send a data bit). Upon receipt of the start bit, the two DTEs are synchronized for the length of time it takes to send the data byte.

Once a data byte is received, the receiving DTE must process it. Processing the character is not an instantaneous operation; at a minimum, the receiving DTE must convert the incoming serial data byte into a parallel byte for storing (this is known as a serial-to-parallel conversion). If the transmitting DTE continues to transmit while the serial-to-parallel conversion is taking place, one or more bytes may be lost unless the receiving DTE is given time to process each one. To give the receiving DTE the necessary processing time, one or more stop bits are appended to the end of each data byte.

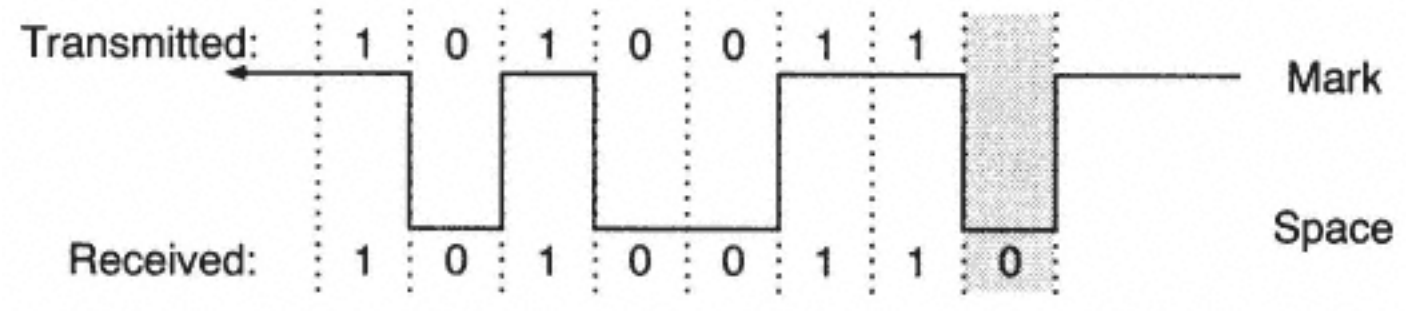

Figure 3.4
Start of transmission is denoted by
the initial mark-to-space transition.

The stop bit(s) are simply the length of time the transmitting DTE must wait before the next byte can be transmitted. The three most common stop bit times are 1, $1\frac{1}{2}$, and 2 bits. Figure 3.5 shows the transmission of the byte *S* with a start bit and two stop bits. The start and stop bits are said to frame the transmitted data byte.

3.2.1
Errors

On occasion, one or more of the bits in the transmission may be changed from a `1` to a `0` (or vice versa) because of noise (typically some form of electromagnetic interference; see Chapter 1) or other interference on the channel. In Figure 3.6, the byte *S* is changed into the byte *Q* because of noise on the line (the second bit is inverted from a `1` to a `0`).

Even with the start and stop framing bits, the receiving DTE has no way of determining from the information received that the byte is correct. What is needed is a mechanism whereby the receiving DTE can detect that an error has occurred in the transmission of the information.

One possible error detection technique is to transmit each byte *twice* and to have the receiver examine both bytes. If the bytes are different, the receiver has detected an error (but cannot determine which byte is the correct one). Interestingly enough, if both bytes are the same, the receiver still has no indication of whether the bytes are correct, since the same error might have occurred to both bytes, resulting in the same bit pattern.

Retransmitting the entire byte presents another problem: Using the channel to send each byte twice means that only half as much information is being transmitted. A much less costly form of error detection can be achieved by

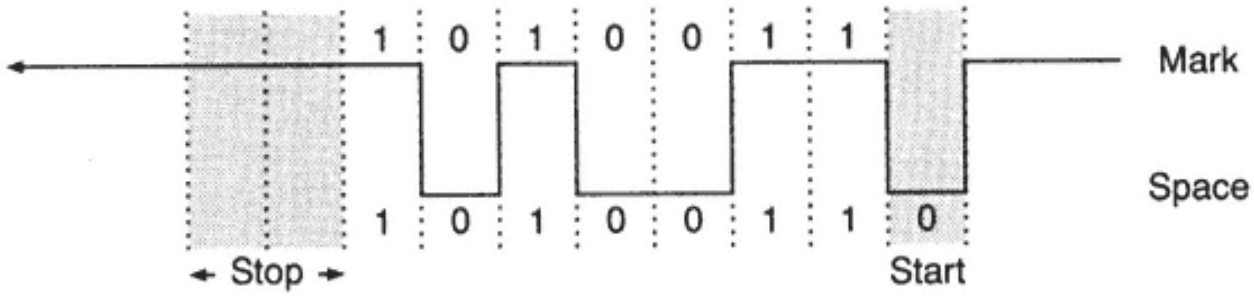

Figure 3.5
The byte *S* sent with a start bit and two stop bits.

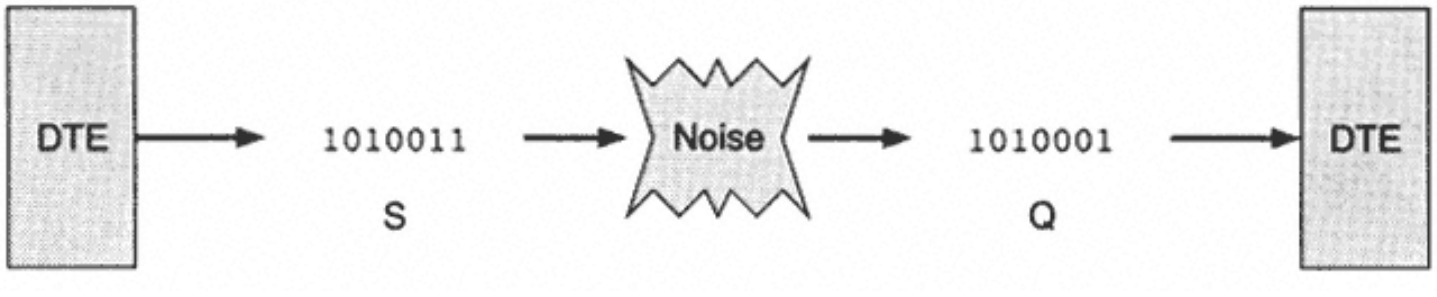

Figure 3.6
Bit inversion due to noise.

summing the bits and determining whether the sum is odd or even. The result of the sum (i.e., odd or even) can be represented as a single bit and transmitted within the frame. This extra bit is known as the parity bit and is transmitted after the last data bit and before the stop bit(s). In Figure 3.7, the byte S is transmitted with a start bit, a parity bit, and two stop bits.

The two best known types of parity are

Odd parity. The sum of all the bits, including the parity bit, is odd. As an example, in Figure 3.8, the byte *S* (bit pattern `1010011`) is transmitted and the parity bit is set; that is, it has a value of `1`, since the sum of the data bits is even (there are *four* data bits set to `1`), but the total of all the bits must be odd. However, as shown in Figure 3.9, if the byte *L* (bit pattern 1001100) is transmitted, the parity bit is cleared; that is, it has a value of 0, since the sum of the data bits is already odd (there are *three* data bits set to `1`).

Even parity. The sum of all the bits, including the parity bit, is even. For example, in Figure 3.10 the byte *S* is transmitted, requiring the parity

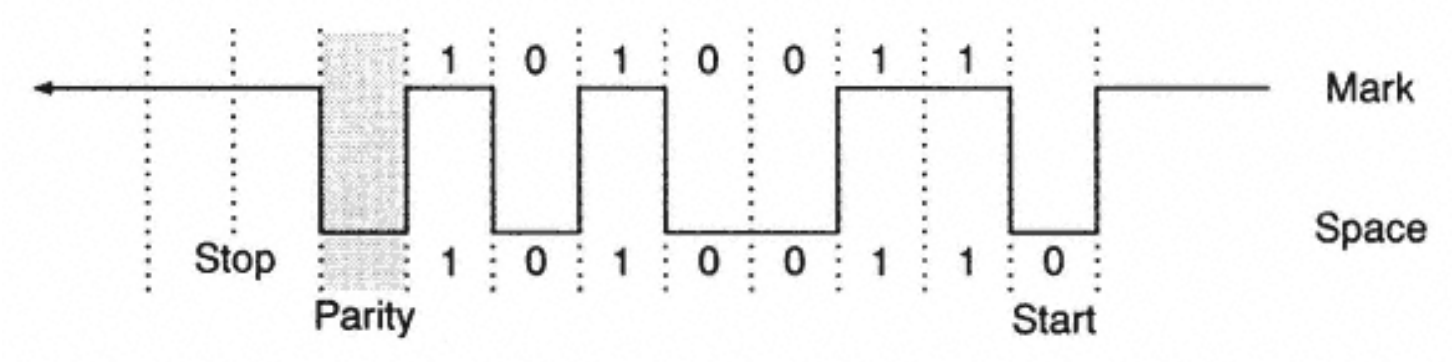

Figure 3.7
The parity bit follows the last data bit and precedes the stop bit(s).

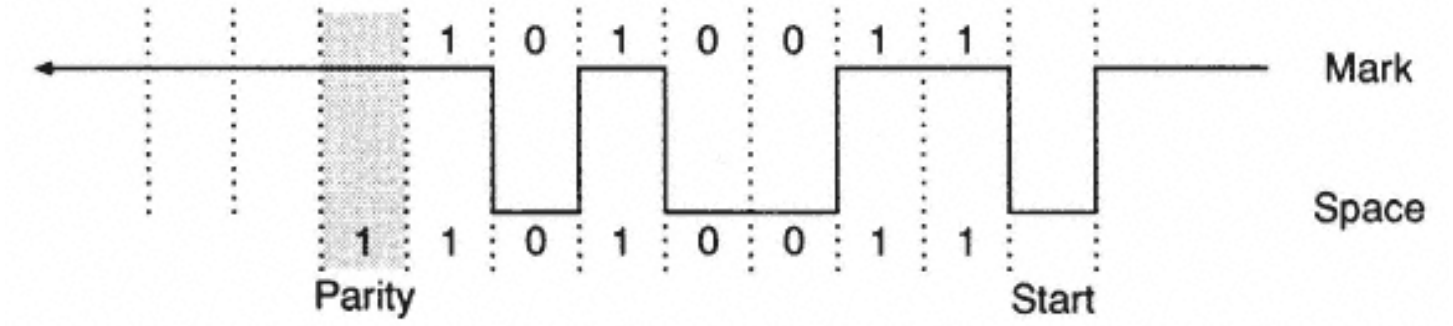

Figure 3.8
Parity bit is set, since the sum of the data bits is not odd.

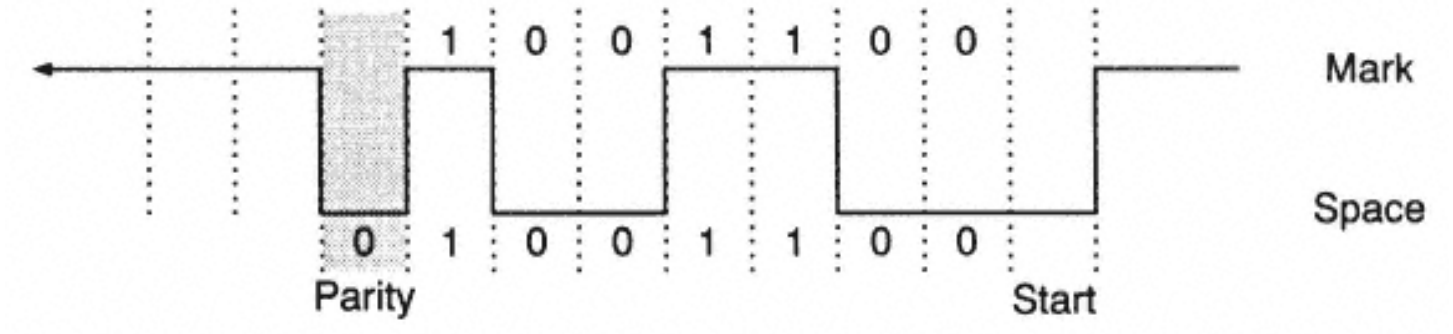

Figure 3.9
Parity bit is cleared, since the sum of the data bits is already odd.

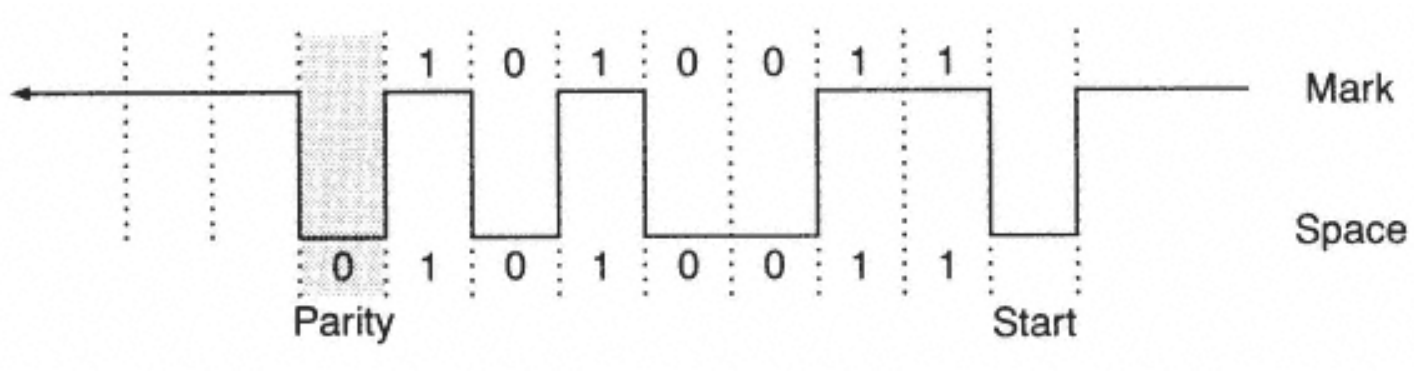

Figure 3.10
Parity bit is cleared, since the sum of the data bits is already even.

bit to have a value of `0` since the sum of the data bits is already even. If the byte *L* (bit pattern `1001100`) is transmitted, the parity bit will have a value of `1`, since the sum of the data bits is odd, but as shown in Figure 3.11, the total of all bits must be even.

The receiving DTE can detect an error when the sum of the bits (data plus parity) is wrong (i.e., odd instead of even or even instead of odd). This type of error is known as a parity error. Note once again that the error has been detected but has not been corrected. If it is to be corrected, the receiving DTE must inform the transmitting DTE that the last byte received was in error.

If the receiving DTE is to detect the parity error, it must support the same parity the transmitting DTE is using (i.e., both must be odd or both must be even). Failure to observe this rule will result in all bytes being flagged as having parity errors. As an example, consider the transmission of the byte *S* with even parity, where there are a total of four bits set to 1 (the parity bit is therefore 0). If the transmitted byte is received by a DTE expecting odd parity, the byte will be received in error because an even number of bits is set.

Now, if interference occurs on the connection between the two DTEs, the receiving DTE can *detect* the error (however, the error cannot be *corrected* since it is not possible to determine which bit(s) is invalid). In Figure 3.12, the byte *S* is transmitted with odd parity, but noise changes the second bit from a `1` to a `0`. Upon receipt of the byte *Q*, the receiver can immediately determine that the data was not received correctly because the parity is wrong (the sum of the bits is *even* when it should be *odd*).

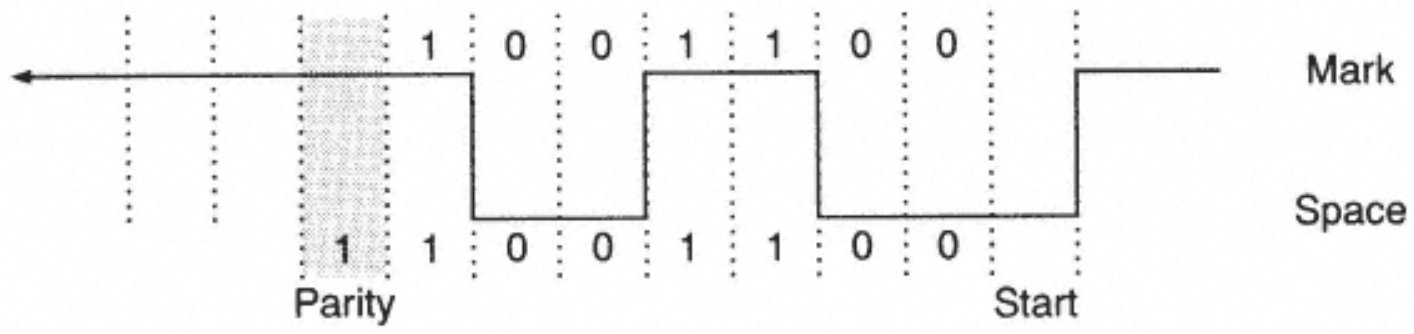

Figure 3.11
Parity bit is set, since the sum of the data bits is not even.

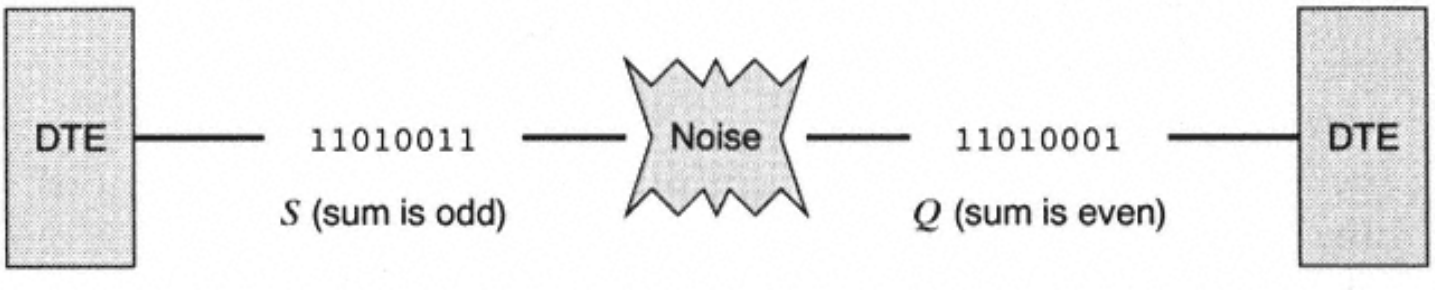

Figure 3.12
Error detection using the parity bit

Does parity error detection still work if the parity bit is changed? Yes, since it works by summing all the bits, data *and* parity. For example, in Figure 3.13, although the byte *S* that is received is valid, the parity is incorrect, indicating that an error has occurred somewhere in the frame. Since there is no indication as to where the error occurred, the entire frame is taken to be invalid.

Parity error detection does not work in all cases, as illustrated in Figure 3.14, in which the byte S is affected by noise that inverts bits 2 and 3, producing the byte *U* (bit pattern `1010101`); note that odd parity is being used. The sum of the bits is *still* odd, which means that the error goes undetected since the parity is correct. Parity checking (odd or even) is not foolproof: it can only detect *odd* numbers of errors.

In addition to odd and even parity, many asynchronous systems also support the following parities (in all cases, the start and stop bits are still sent):

None. The parity bit is neither generated nor transmitted. This reduces the number of bits in the frame, which can be useful when transferring large

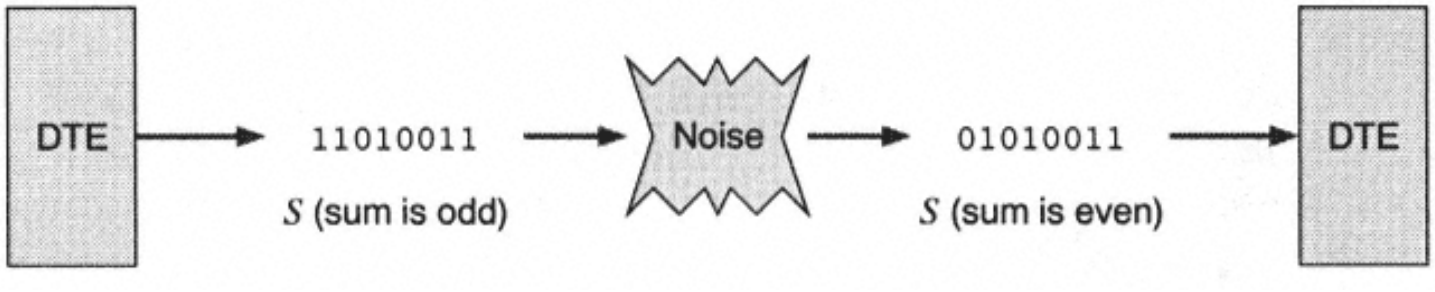

Figure 3.13
Error detection using the parity bit (parity bit inverted).

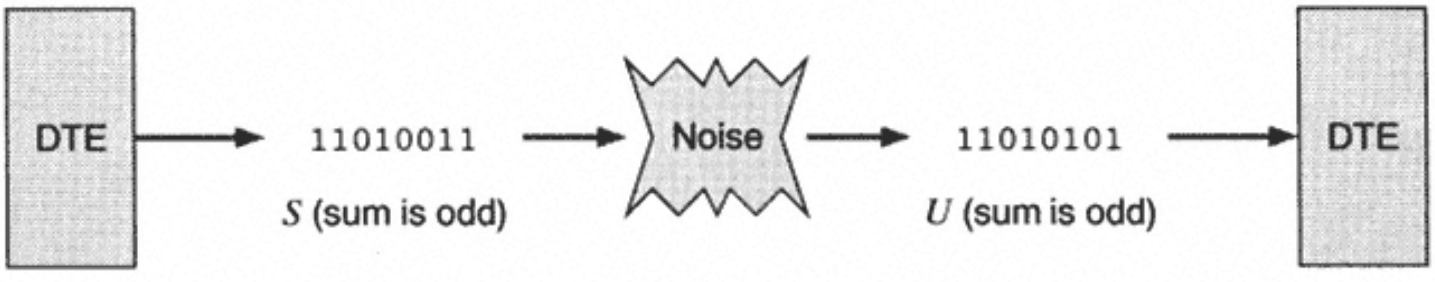

Figure 3.14
An example of parity failing to detect an error.

amounts of data (see Chapter 6). However, it also means that the receiver cannot detect if an error has occurred in the transmission of the data.

Mark parity. The parity bit, regardless of the data bits, always has a value of `1` (hence the name "mark parity"). Figure 3.15 illustrates the transmission of the byte *S* with mark parity. If the parity bit received is *not* a mark, then a parity error is detected; otherwise, the byte is assumed to be error-free.

Space parity. The parity bit, regardless of the data bits, always has a value of `0` (hence the name "space parity"). The transmission of the byte *L* with space parity is illustrated in Figure 3.16. If the parity bit received is *not* a space, then a parity error is detected; otherwise, the byte is assumed to be error-free.

3.2.2
Other Errors

There are two other common errors that can occur on an asynchronous channel. An overrun error occurs when the receiving DTE does not process the incoming data bytes quickly enough. In Figure 3.17, three bytes have been sent by the transmitting DTE; however, before the receiving DTE has had time to receive and process the second byte (it is still processing the first byte), the third byte arrives, overwriting the second byte. Overruns can be avoided by reducing the processing load on the receiving DTE or by reducing the line speed.

A framing error occurs when the receiving DTE cannot detect the end of the frame (i.e., the stop bits cannot be detected). There are two reasons for the DTE missing the stop bit(s). One, noise may have caused the stop bits

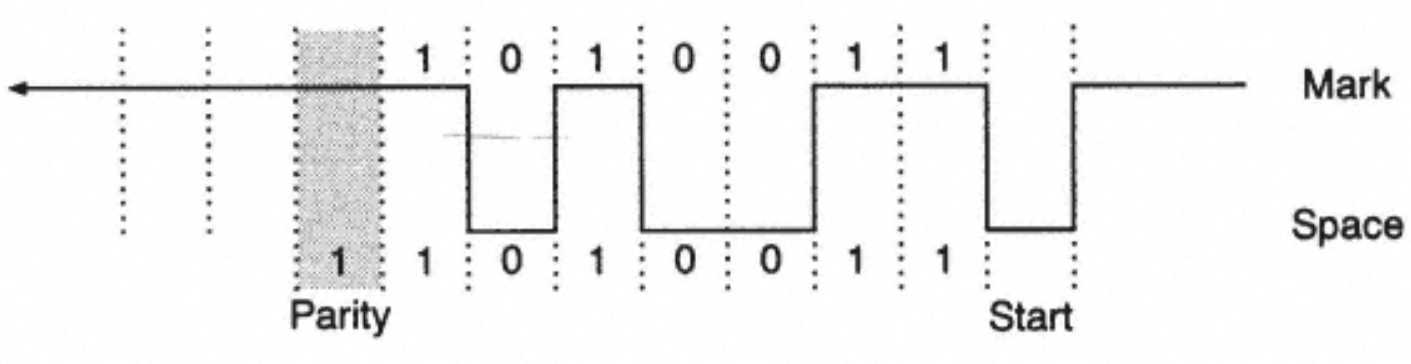

Figure 3.15
In mark parity, the parity bit is always set.

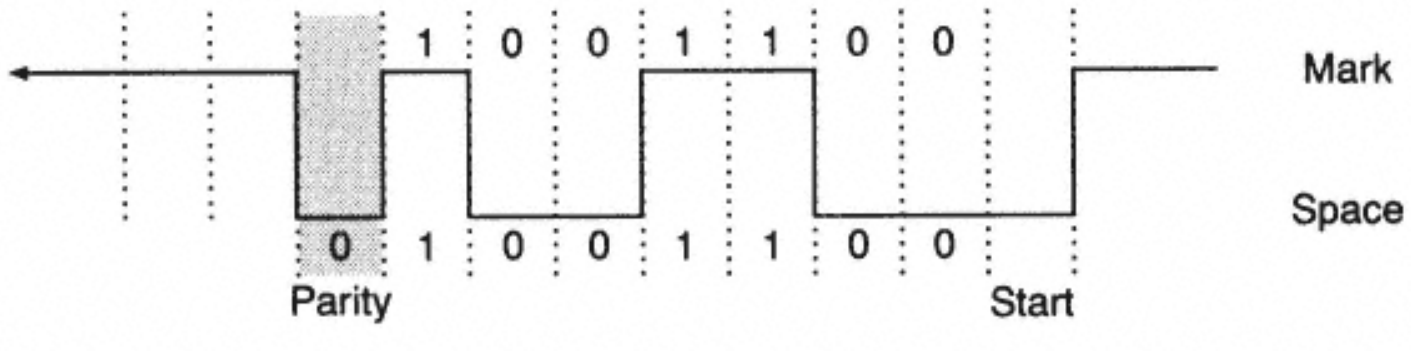

Figure 3.16
In space parity, the parity bit is always cleared.

Transmitting DTE		Receiving DTE
First byte sent	⟶	Receives first byte
Second byte sent	⟶	Processing first byte
Third byte sent	⟶	Receives third byte

Figure 3.17
An example of an overrun error.

to be inverted, prompting the receiving DTE to detect one or more spaces instead of marks. Two, noise may have caused the start bit to be inverted, causing the DTE to start assembling the byte with the wrong bit. Regardless of the cause of the framing error, there is no mechanism available to allow the receiving DTE to determine the correct value of the transmitted byte (other than the retransmission of the byte).

Further explanation of both overrun and framing errors is given in Section 3.3.

3.2.3
The Break Condition

In addition to the characters found on a terminal's keyboard, many DTEs with keyboards have a key marked Break. The Break key does not have an equivalent ASCII (or EBCDIC) character code; rather, it is intended to cause the user's DTE to generate a break condition on the channel. A break can be used for a number of reasons; in many applications it is intended to signal the receiving DTE that attention is required.

As shown in Figure 3.18, a break condition is indicated by holding the channel in the space state for *at least* the amount of time required to send an entire frame (i.e., a start bit, the data bits, the parity bit, and the stop bits).

How the receiving DTE handles the break condition is determined by a higher level of protocol. For example, software may be designed to treat the reception of a break as an indication that a process being executed is to be aborted.

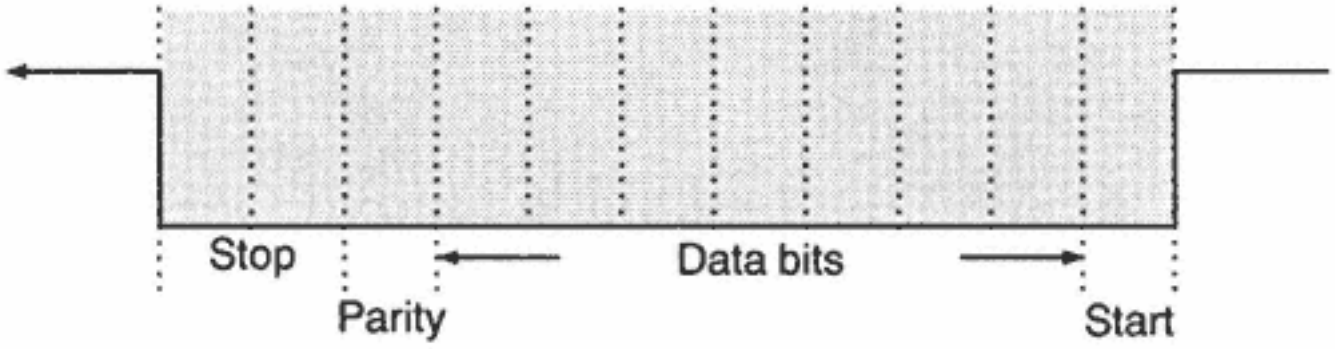

Figure 3.18
In a break condition, all the frame bits are cleared.

3.3
Asynchronous Communications Hardware

Data within a DTE is usually stored and moved in a parallel fashion (the start and stop framing bits do not exist); however, data sent across the channel is generally sent serially. The DTE is therefore required to convert the internal data to a form that can be transmitted on the channel (in serial form, with start, stop, and parity bits). Similarly, the serial data received from the channel must be changed into parallel for use within the DTE.

The conversion of the data for transmission and reception is done in hardware by a device known as a Universal Asynchronous Receiver/Transmitter, or UART. Figure 3.19 represents the major UART connections.

The UART is responsible for generating the start, stop, and parity bits, as well as for removing them. Channel conditions such as the break condition, framing errors, and overrun errors are signaled by the UART back to the processor. The processor can send control signals to the UART, including line speed, word size, parity, and the number of stop bits.

3.3.1
UART Transmission

When transmitting data, the UART must convert the internal parallel byte into a stream of serial bits using parallel-to-serial conversion. The byte to be transmitted is supplied to the UART by the DTE's processor. The bits are stored in a shift register (often referred to as the Transmit Shift Register). The start bit is sent first, followed by the bits in the shift register. As shown in Figure 3.20, the shift register clocks them onto the channel, one at a time, at the line speed.

Table 3.1 shows how and when the parity bit is generated (and subsequently sent), depending upon the parity selected. Odd and even parities are calculated by the exclusive or'ing of the data bits as they are shifted onto the channel. The parity bit is initialized to `0` for even parity and `1` for odd parity. Figure 3.21, illustrates how a feedback loop can be used to calculate

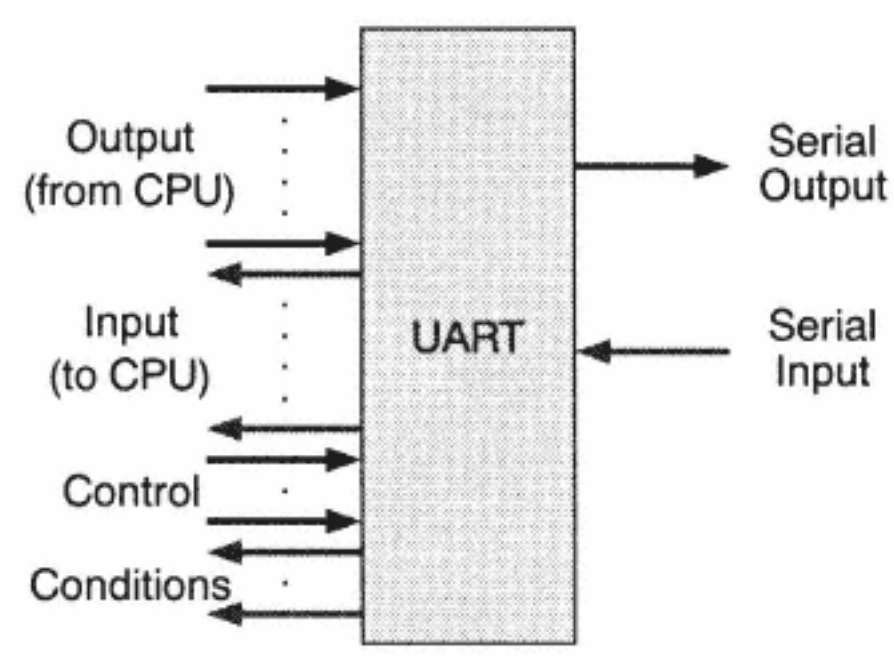

Figure 3.19
The basic UART connections.

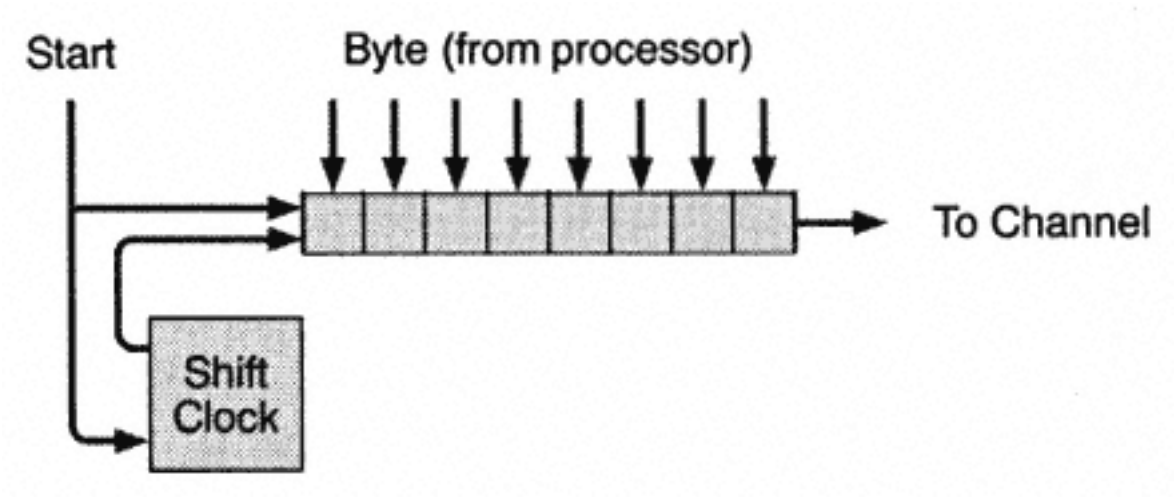

Figure 3.20
The UART shift register.

Table 3.1 UART parity bit generation

Parity	How Generated
None	Parity bit is not sent
Odd	Generated by exclusive or'ing the outgoing data bits
Even	Generated by exclusive or'ing the outgoing data bits
Mark	Parity bit is set to `1`
Space	Parity bit is set to `0`

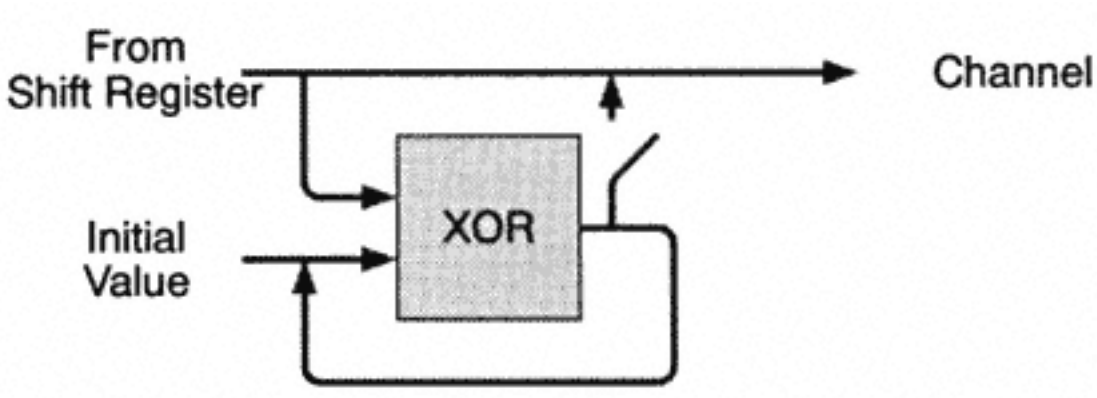

Figure 3.21
A feedback loop for calculating parity.

the parity (the initial value of the parity bit depends upon the parity being calculated).

Each bit that is sent is included in the parity generation. Once the final data bit is included in the parity, the parity bit can be clocked onto the channel for transmission. If mark or space parity is selected, a `1` or `0`, respectively, is put onto the channel after the final bit, regardless of the value in the parity generation hardware.

The parity bit is followed by one or more stop bits. Stop bits are "generated" by the UART entering an idle state, in which the channel idles with marks and blocks transmission for at least the time it would take to transmit the required number of stop bits.

It is important to note that the line speed indicates the *total* number of bits the UART can transmit in one second. Therefore, when determining

the number of bytes that can be transmitted in one second, it is necessary to include the number of data bits as well as the number of start, stop, and parity bits. For example, a 1200-bps line transmitting 7-bit bytes, with parity (odd or even), one start and one stop bit (a total of 10 bits) could transmit a maximum of 120 bytes per second.

Once the UART has sent a byte, the processor can be signaled that the UART is ready to transmit another byte.

If data is supplied by the processor to the UART only when the shift register is empty, noticeable transmission delays may result. This is because transferring a byte to the UART is not an instantaneous operation since the line speed is considerably less than the speed of the processor and the UART. Figure 3.22 illustrates this problem.

The periods when the channel is idle can be reduced if a second register, a holding register (usually referred to as the Transmit Holding Register), is placed between the processor and the UART. The relationship between the two registers is shown in Figure 3.23.

Time	Processor	UART	Channel
1	Supplies byte to UART		Idle
2		Shifts	Active
3		Data	Active
4		Bits	Active
5		Signals Completion to Processor	Idle
6	Supplies byte to UART		Idle
7		Shifts	Active
8		Data	Active
9		Bits	Active
10		Signals Completion to Processor	Idle

Figure 3.22
Throughput reduction due to speed of UART.

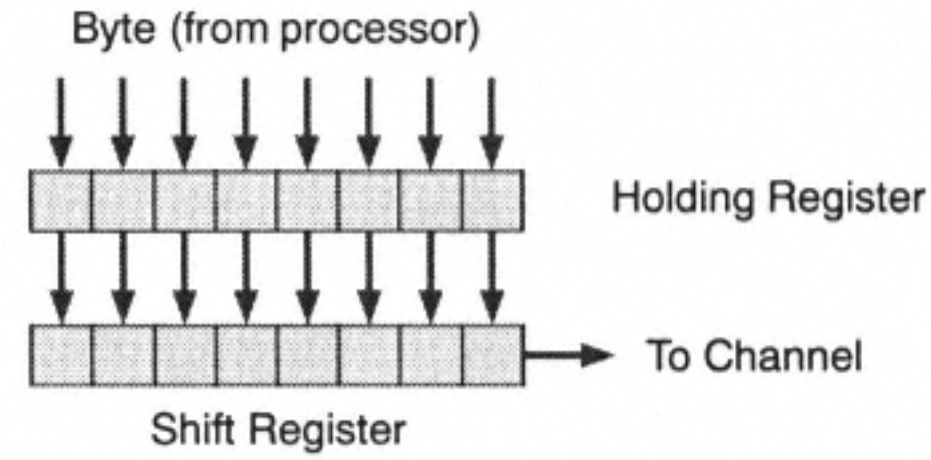

Figure 3.23
UART transmit holding register.

Time	Processor	Holding Register	Shift Register	Channel
1	Supplies byte to HR	Empty	Empty	Idle
2		Supplies byte to SR	Empty	Idle
3		Signals processor	Shift	Active
4	Supplier byte to HR	Empty	Data	Active
5		Supplies byte to SR	Bits	Active
6		Signals processor	Shift	Active
7	Supplier byte to HR	Empty	Data	Active
8		Supplies byte to SR	Bits	Active

Note: "HR" denotes holding register to "SR" denotes shift register.

Figure 3.24
How channel throughput is increased using a holding register.

As shown in Figure 3.24, delays are now reduced because the UART can be transmitting a byte from the shift register while the processor fills the holding register. This technique, whereby two registers are used in a transmission, is known as double buffering. In addition to accepting bytes from the processor, most UARTs can signal the processor when the Transmit Holding Register is empty and when the Transmit Shift Register is empty.

3.3.2 UART Reception

The UART is in an idle receive state while the channel is in an idle (mark) state. Upon detection of a channel state change (i.e., the channel changes from a mark to a space), the UART determines the center of the start bit and then reads bits from the channel at regular intervals using a sampling clock. The bits are assembled in a shift register; each time the sampling clock "ticks," the register is shifted to the right and the next bit from the channel is read. When all of the bits have been read from the channel (the number of bits read is determined by the word size), the resulting byte is supplied to the processor (see Figure 3.25).

The UART also scans for the start and stop bits and, if required, calculates and performs a parity check. How parity is handled depends upon the parity chosen, as the following paragraphs explain.

None. The UART ignores parity checking altogether.

Mark or space. The channel can be sampled after the last data bit is read. The value of the bit can be compared to that expected by the receiving UART. If the parity bit received is the same as that expected, the data is assumed to be correct; otherwise, a parity error is signaled.

Odd or even. Both the incoming data bits *and* the parity bit are used to determine whether a parity error has occurred. The UART initializes its parity check circuit to 0 (even parity) or 1 (odd parity) and proceeds to

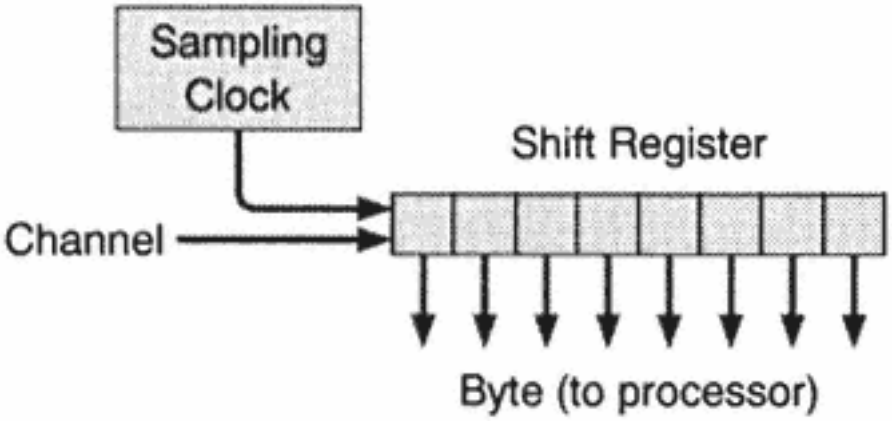

Figure 3.25
UART reception shift register.

generate parity by the exclusive or'ing of the incoming data bits with the parity. If no errors are detected, the value of the parity calculated by the UART should be the same as that received by it. At this point, the UART can explicitly compare the parity bit received with that calculated and signal the DTE accordingly.

Alternatively, the UART can include the incoming parity bit in the parity calculation. If the parity in the frame received is the same as that calculated by the UART (i.e., both 0 or both 1), the result of this final exclusive or'ing is 0. However, if the values are different, the final exclusive or'ing will result in a value of 1. This result can be made available to the DTE; for example, a 0 could indicate that no errors were detected, while a 1 would signal a parity error.

Once a complete byte has been assembled in the shift register, the UART can signal the processor that a byte is available. At this point, the processor can copy the byte from the shift register into a memory location.

If the assembled byte is kept in the shift register and the processor does not read the shift register before the first bit of the next byte arrives, an overrun error will occur. Since there is only one shift register, either the incoming byte must be discarded or the byte in the shift register must be overwritten. Ideally, no bytes will be lost (neither the incoming byte nor the byte in the shift register), meaning that the byte in the shift register should be read by the processor as quickly as possible. The reception technique is similar to that used in transmission: a second register is placed between the processor and the shift register (i.e., double buffering). As soon as a byte is assembled in the shift register, it is copied (in parallel) into the holding register (see Figure 3.26).

The addition of the holding register does not eliminate the possibility of an overrun occurring; it merely reduces the probability. For example, a byte could be copied into the holding register and the processor signaled; if the processor did not remove the byte from the holding register before a second byte was copied into it, an overrun error would still occur.

The other errors discussed in Section 3.2.1 are also handled by the UART's reception hardware. If the parity calculated as the byte is received

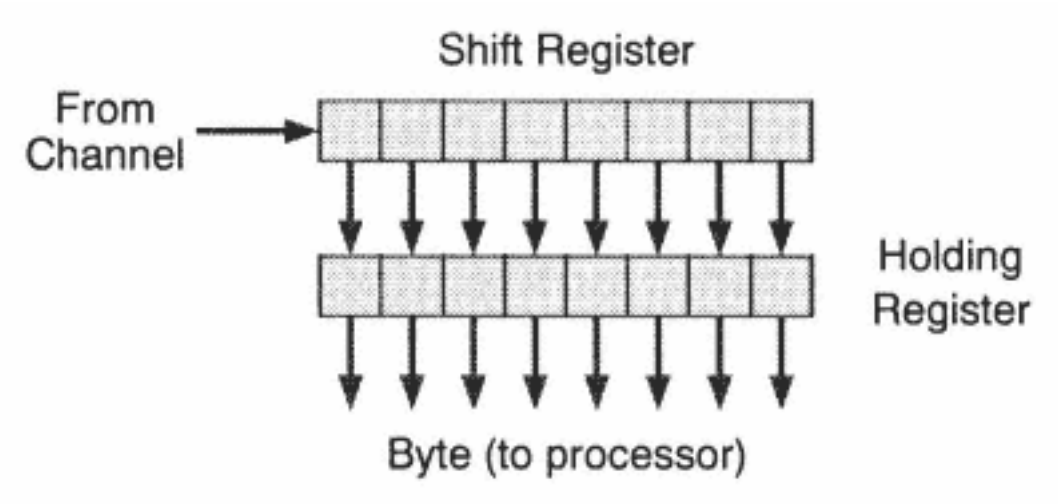

Figure 3.26
UART reception holding register.

does not agree with the parity supplied with the byte, the UART should signal the processor that a parity error has occurred. If a stop bit (i.e., a mark) is not detected after the last bit (data or parity, depending on whether parity is required), the UART should signal the processor that a framing error has occurred.

The UART should also be able to detect a break condition, specifically after the line has been in the space state for at least the start bit, data bits, parity bit, and stop bits. Often, a UART will signal a framing error when detecting a break because the framing error condition can be indistinguishable from the break condition. For example, an inverted (missing) stop bit on a null character sent with even parity is equivalent to a break condition.

The value of the byte assembled by the UART clearly depends upon the bits that are sampled from the channel. Although the signal diagrams have shown the rising and falling edges of a transmitted character to be sharp and square, in reality the edges may be much less distinct. In fact, over time the signal may become attenuated or weakened, so that it is extremely difficult to determine the value of the bit, as shown in Figure 3.27.

Since the value of a bit may become harder to determine at its edges, UARTs are normally designed to sample the line as close to the *center* of each bit as possible. The center of a bit is found by first locating the center of the start bit. Once the center of the start bit has been found, the UART can sample the line once per bit time (i.e., the inverse of the line speed; for example, a 1200-bps line would have a bit time of $\frac{1}{1200}$ of a second) to find the center of the remaining bits in the transmission.

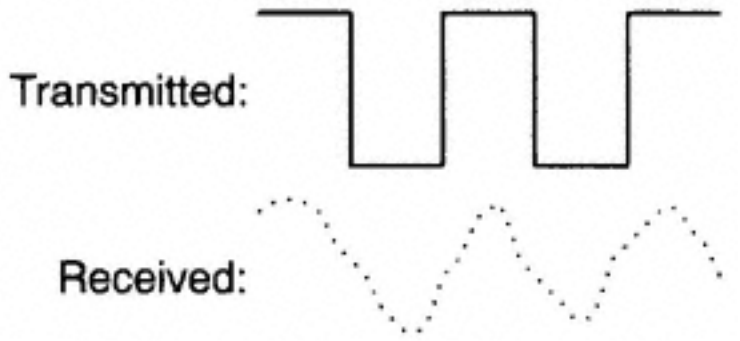

Figure 3.27
Signal attenuation.

The algorithm to find the center of the start bit is as follows:

1. Wait for a mark-to-space transition.

2. When a transition has been found, sample the line at 16 times the bit rate.

3. After 8 "ticks" (i.e., one-half a bit time), sample the line:

    If a space (0) is found, a start bit is assumed and the line can be sampled once per bit time.

    If a mark (1) is found, a false start bit was found, and the line sampling can be abandoned.

For example, if the letter *S* is sent, the sampling of the channel occurs as shown in Figure 3.28.

Since the UART may not start its sampling clock at the exact edge of the start bit's mark-to-space transition, using a 16-times counter may prove to be inaccurate in determining the bit's center. The error can be made smaller by sampling at 32 or even 64 times the bit rate.

The underlying assumption in sampling the channel is that both the transmitter and receiver agree to a common bit rate (or line speed) that does not vary over the period of the transmission of the frame. However, should one of the clocks be slower than the other, clock drift may occur, possibly resulting in the misinterpretation of the byte or an overrun error. For example, in Figure 3.29, the receiver samples the line at a rate faster than the transmitter sends (the byte *S* is transmitted with 7-bit, even parity).

Because the receiver is sampling the line at a rate faster than the transmitter sends, the value of the byte received by the DTE will not be the same as was transmitted. Depending upon the degree of signal attenuation, the DTE could receive one of a number of possible byte values, for example (the leftmost bit is the parity bit):

`10100111`. The third bit received (sampled at the space/mark transition between the second and third bits transmitted) is read as a `1`. A parity error will occur.

`10100011`. The third bit is read as a `0`, meaning that the parity bit and the data bits are assumed to be correct.

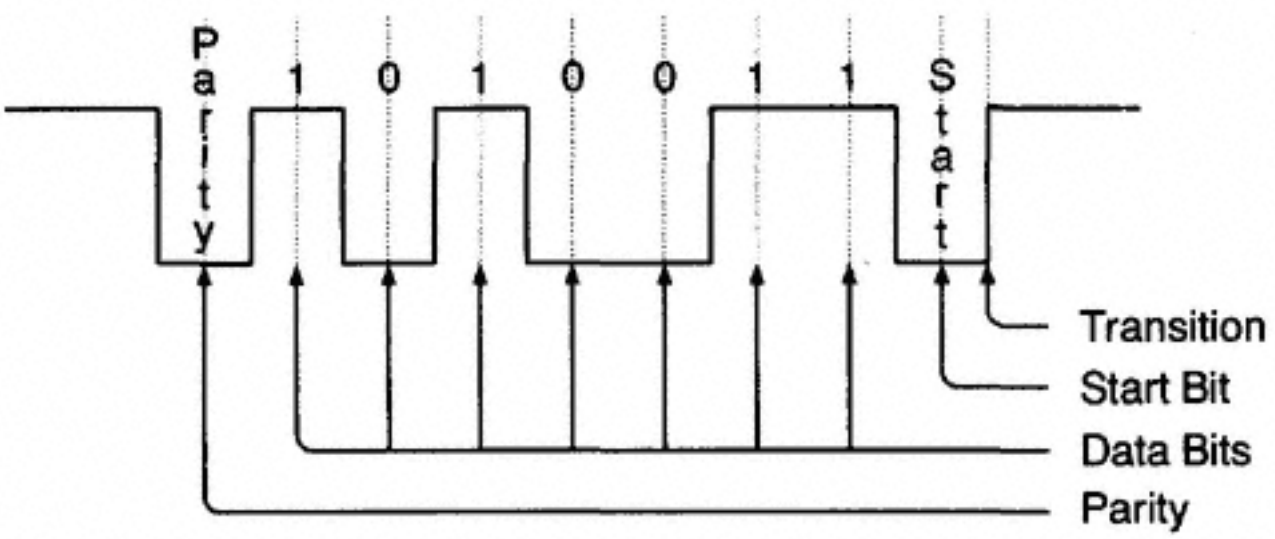

Figure 3.28
The channel is sampled at the center of each bit.

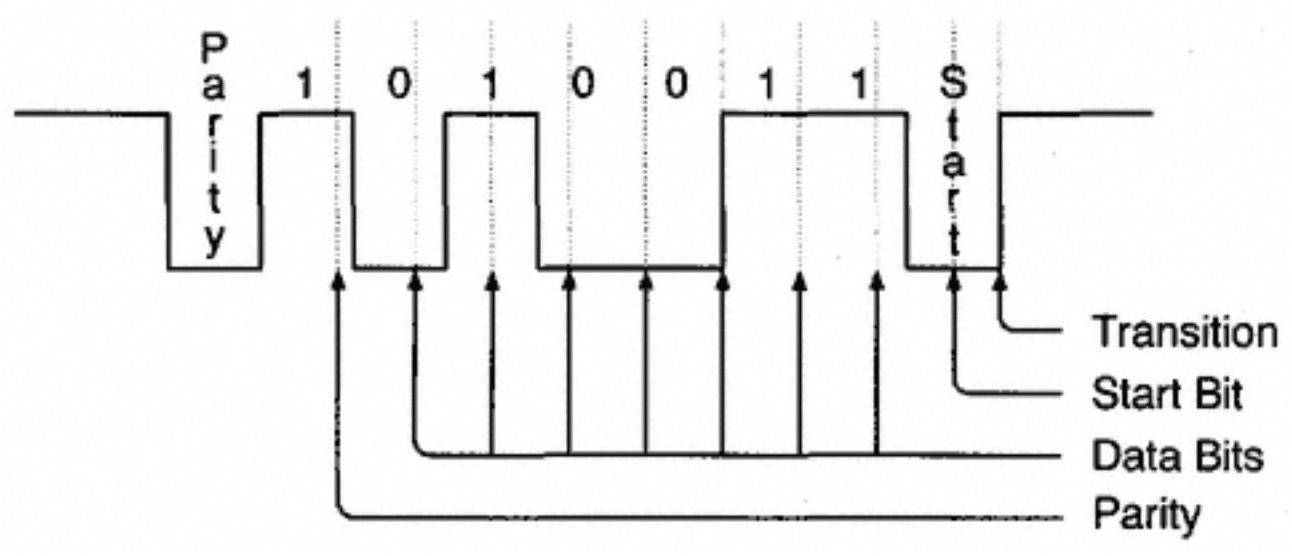

Figure 3.29
Clock drift: the center of the bit is not found.

In each of the above cases, the parity bit may be misinterpreted as the start bit of the next byte. Regardless of the value of the byte received by the DTE, it is wrong, meaning that one or both UART clocks need to be reinitialized and the data retransmitted.

3.4 The PC UART

A typical PC is supplied with at least one UART (also referred to as a serial interface, asynchronous communication adapter, or serial port). Although there is no single supplier of UARTs, most PC UARTs are compatible with the Intel 8250 UART. Commkit supports the 8250 UART.

The 8250 UART is programmable and permits

A variable word size (5-, 6-, 7-, or 8-bit words).

A variable number of stop bits (1, $1\frac{1}{2}$, or 2).

A choice of parity (none, mark, space, odd, or even).

A choice of line speeds (50 to 9600 bps or higher).

The UART permits the transmission and reception of data and can signal the processor with the following status indications:

Data ready (i.e., data is available in the receive buffer).

Reception errors (i.e., overrun, parity, and framing).

Break condition detected.

Transmit shift register empty.

Transmit holding register empty.

The UART also supports interrupts, three of which are of interest at this point:

Receive line status (a change in the status of the line has occurred, typically an error or a break).

Data available (a byte has been received).

Transmit holding register empty (a byte has been copied into the shift register and the UART can accept another byte).

UART interrupts occur through device vectors 11 (serial port 2) and 12 (serial port 1). The type of interrupt can be determined by the processor reading the Interrupt Identification Register (see below).

3.4.1
Programming the UART

As with the other PC devices that have been discussed, the UART can be accessed through a series of ports using the three instructions: `inportb()`, `outportb()`, and `outport()`.

Although there are a total of seven ports for the UART, only five will be discussed in this chapter (the remaining two are discussed in Chapter 8). The five ports, their addresses, offsets, and uses are shown in Table 3.2 (note that some ports have multiple uses).

For the sake of brevity, the remainder of this section will refer only to serial port 1. However, everything discussed for serial port 1 is applicable to serial port 2.

UART Initialization

When a computer is first powered on, the hardware is often in an undetermined state; for example, the contents of the computer's memory are unknown. Similarly, the internal settings of the UART are unknown and must be initialized. That is, the line speed, word size, parity, and number of stop bits must be written to the UART before it is used. Three registers serve to initialize the UART: the two line speed registers (`0x3F8` and `0x3F9`) and the Line Control Register (`0x3FB`).

The UART's line speed is generated by dividing its internal clock rate (1.8432 MHz) by a programmable 16-bit divisor (set up in the two line speed registers, `0x3F8` and `0x3F9`). The value resulting from the division is *16 times* the actual line speed. Once a mark-to-space transition has been

Table 3.2 UART Port Addresses and Functions.

Port 1	Port 2	Address Offset	Uses
`0x3F8`	`0x2F8`	Base address* + 0	Transmission Register Buffer Receive Register Buffer Line Speed (Least Significant Byte)
`0x3F9`	`0x2F9`	Base address + 1	Interrupt Enable Line Speed (Most Significant Byte)
`0x3FA`	`0x2FA`	Base address + 2	Interrupt Identification Register
`0x3FB`	`0x2FB`	Base address + 3	Line Control Register
`0x3FD`	`0x2FD`	Base address + 5	Line Status Register

*The base address refers to `0x3F8` (Port 1) or `0x2F8` (Port 2).

detected, the center of the start bit can be determined by examining the channel after eight "ticks" of the clock have occurred. As soon as the center of the start bit has been found, the center (and value) of each subsequent bit can be obtained after every sixteen "ticks" of the clock. The value selected for the line speed is used for both transmission and reception.

Although any 16-bit value between 1 and 2161 can serve as the divisor, specific values must be used in order to obtain "well-known" line speeds. Table 3.3 lists some of the divisors and their associated line speeds.

As an example, consider the effect of setting the line speed to 9600 bits per second. First, the UART's internal clock rate of 1.8432 MHz is divided by `0x0C` (decimal 12) to obtain the line speed: 0.1536 MHz or 153,600 ticks per second (remember, this is still 16 times the actual line speed). Every 16 of these ticks is equivalent to the time taken to send a single bit, resulting in an effective speed of 153,600 divided by 16, or 9600 bits per second. (To find the center of the start bit, the UART samples the channel 8 ticks after the mark-to-space transition is found.)

The 16-bit value of the line speed divisor is written into the two line speed register ports. Since the two line speed registers (`0x3F8` and `0x3F9`) are used both in the transmission and reception of data as well as to indicate which interrupt has occurred, the UART must be informed of the function of the two registers *before* they are used. If bit 7 of the Line Control Register (the Divisor Latch Access Bit, or DLAB) is set, ports `0x3F8` and `0x3F9` can be accessed as the line speed divisor registers (see Figure 3.30).

The line speed register ports can now have the value of the line speed

Table 3.3 Some lines speeds and their associated divisors.

Divisor (Hexadecimal)	Line Speed
`0x0900`	50
`0x0180`	300
`0x0060`	1200
`0x0030`	2400
`0x000C`	9600

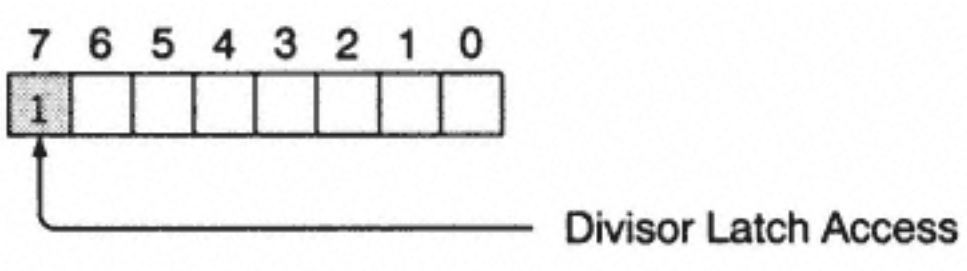

Figure 3.30
Line Control Register: setting the Divisor Latch Access Bit.

written into them: the lower 8 bits into `0x3F9` and the upper 8 bits into `0x3F8`. In Figure 3.31, the line speed is to be set to 9600 bps, meaning that register 0x3F8 will be set to 0x00 and register `0x3F9` to `0x0C`.

The word size, parity, and stop bits are also initialized through the Line Control Register, the contents of which are shown in Figure 3.32 (note that the Divisor Latch Access Bit is cleared). The terms used in Figure 3.32 are defined as follows:

Word length. The word length is specified in bits 0 and 1 of the Line Control Register, permitting four different word lengths (see Table 3.4).

Number of stop bits. Bit 2 of the Line Control Register signifies the number of stop bits in each transmitted or received frame. The number of stop bits generated depends upon whether this bit is set and on the number of bits in each word (see Table 3.5).

Parity enable. When the parity enable bit (bit 3) is set, the UART both generates and checks frame parity (the type of parity, even or odd, is specified in bit 4).

Even parity select. Bit 4, even parity select, is set to indicate even parity and cleared to indicate odd parity. This bit has meaning only if bit 3 (parity enable) is set.

Stick parity. Stick parity is a mechanism whereby the transmitted byte can be sent (and received) with either mark or space parity. When both the stick parity bit and the parity enable bits are set, parity generation is disabled and the value of the parity bit is set to the inverse of whatever parity (odd or even) has been selected in bit 4. Table 3.6 shows the required settings

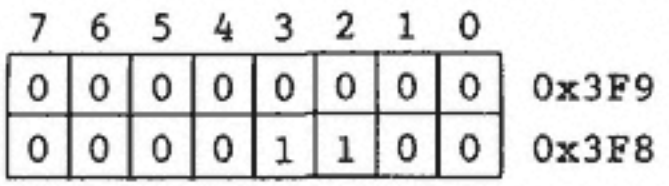

7	6	5	4	3	2	1	0	
0	0	0	0	0	0	0	0	0x3F9
0	0	0	0	1	1	0	0	0x3F8

Figure 3.31
Line Speed Register: setting
the line speed to 9600 bps.

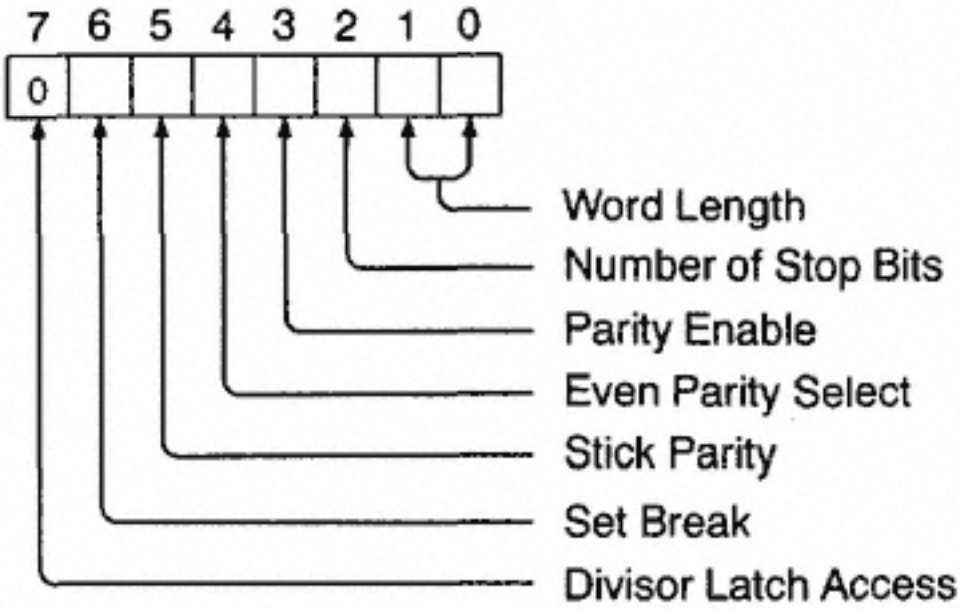

Figure 3.32
Line Control Register.

Table 3.4 Line Control Register: word length bit values.

Bit 1	Bit 0	Word Length (bits)
0	0	5
0	1	6
1	0	7
1	1	8

Table 3.5 Line Control Register: stop bit values.

Bit 2	Word Length	Number of Stop Bits
0	Any length	1
1	5-bit	1 1/2
1	6-, 7-, 8-bit	2

Table 3.6 Line Control Register: stick bit values.

Parity Required	Parity Enable	Even Parity Select
Mark	1	0
Space	0	1

to generate mark or space parity and assumes that the stick parity bit is set.

As always, both the transmitting and receiving DTEs must agree to the same parity (mark, space, odd, even, or none).

Set break. The set break bit can be used to force the line into a break condition. That is, when this bit is set, the line is held in the space state until the bit is cleared. Transmitting a break requires the processor to determine the amount of time needed for a break, given the line speed, and then to set the set break bit. Once the time has expired, the processor can clear the set break bit.

Divisor Latch Access Bit. The Divisor Latch Access Bit (bit 7), when set, allows the processor to change the line speed (see discussion above). When cleared, the speed cannot be altered.

The Line Control Register settings are applicable to both transmission and reception. For example, if the UART is initialized to 7-bit data and even

parity, the UART will transmit *and* receive 7 bits of data per frame with even parity.

To initialize the UART to 7-bit data, 1 stop bit, and even parity, the Line Control Register will contain the information shown in Figure 3.33.

The C software used to set the primary port UART to 9600 bps, 7 bit bytes, 1 stop bit, and even parity could be written as follows:

```
#define DLABon      0x80
#define DLABoff     0x00
#define BPS96K      0x0c
#define SEVENBIT    0x02
#define STOPBITS    0x00
#define PRTYENA     0x08
#define EPS         0x10

#define DIVLTCHLSB 0x3F8
#define DIVLTCHMSB 0x3F9
#define LCR         0x3FB

initialize ()
{
outportb(LCR, DLABon);
outport(DIVLTCHLSB, BPS96K);
outportb(LCR, DLABoff + SEVENBITS + STOPBITS + PRTYENA + EPS);
}
```

UART Transmission

The UART has two registers associated with data transmission: the Transmit Holding Register and the Transmit Shift Register. Although the UART can signal when either of these registers is empty, the processor can write only to the Transmit Holding Register.

The Transmit Holding Register is an 8-bit register accessed through port `0x3F8`. A byte written to this port (using `outportb()`) is normally copied by the UART to the shift register and transmitted serially. For example, to transmit the character A:

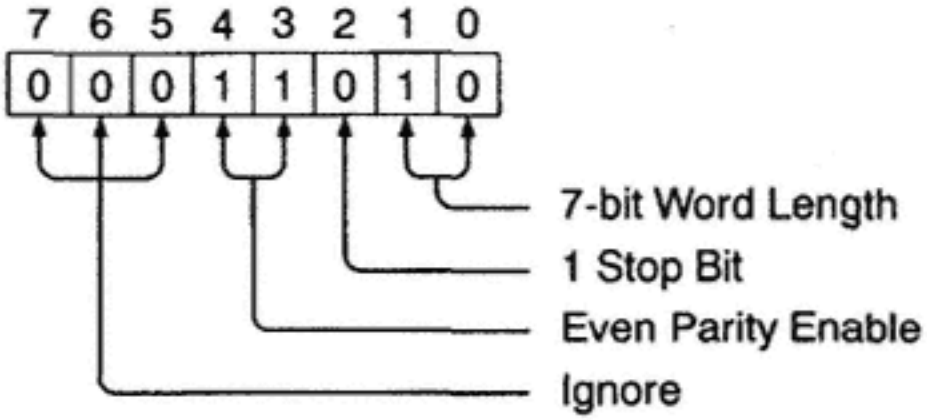

Figure 3.33
Line Control Register settings for
7-bit data, 1 stop bit, even parity.

```
/* UART has been initialized */
outportb(0x3F8, 'A');
```

However, should the byte be written to the port before the previous byte has been copied into the shift register, data will be lost: a form of internal overrun will occur. For example:

```
main()
{
char ch;
/* Initialize UART */
for (ch = 'A'; ch <= 'Z'; ch++)
     outportb(Ox3F8, ch):
}
```

In the above code fragment, the processor supplies bytes to the UART faster than they can be transmitted. Accordingly, not all of the bytes will be sent, since the UART is slower than the processor. This problem can be overcome by reducing the rate at which the processor supplies the data to the UART, either by polling the UART to determine when the next character can be supplied or by allowing the UART to interrupt, signaling that the Transmit Holding Register is empty.

Polling of the UART is done through the Line Status Register, port 0x3FD. The Line Status Register signals the condition of the UART's transmission and reception facilities (see Figure 3.34).

The status of the Transmit Holding Register Empty flag and that of the Transmit Shift Register Empty flag are indicated by their values: 0 (the register contains data) or 1 (the register is empty and data can be written to it). To maximize throughput, data should be written to the Transmit Holding Register as soon as the Transmit Holding Register Empty bit is set. A typical polling sequence could be written as follows:

```
#define TXR       0x3F8
#define LSR       0x3FD
#define THRE      0x20
main()
{
```

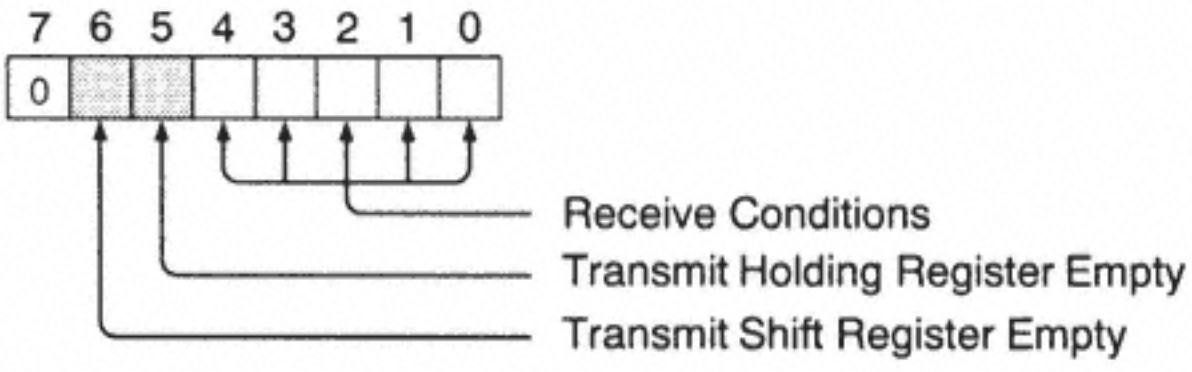

Figure 3.34
Line Status Register: transmission status bits.

```
char ch;
/* Initialize UART */
for (ch = 'A'; ch <= 'Z'; ch++)
{
      while ((inportb(LSR) & THRE) == 0);
      outportb(TXR, ch);
}
}
```

UART transmission using interrupts is described later in this chapter.

UART Reception

The UART makes data available to the processor through the Receive Buffer Register (port `0x3F8`). The port is read and a copy of whatever is in the Receive Buffer is returned to the processor:

```
#define RCVR     0x3F8
main()
{
/* Initialize UART */
for(;;)
     printf(" %c\n", inportb(RCVR));
}
```

As when supplying data to the Transmit Holding Register, the UART assembles characters at a much slower rate than the processor can read them, meaning that the Receive Holding Register should be accessed only after a byte has been read from the line. The processor can determine whether a byte is available by reading the Line Status Register (`0x3FD`) (see Figure 3.35).

The Data Ready bit is used to signal that a byte has been assembled and placed in the Receive Buffer Register (`0x3F8`). A value of 0 means that no new data has been placed in the buffer since the last time it was read, while a value of 1 indicates that new data is now available.

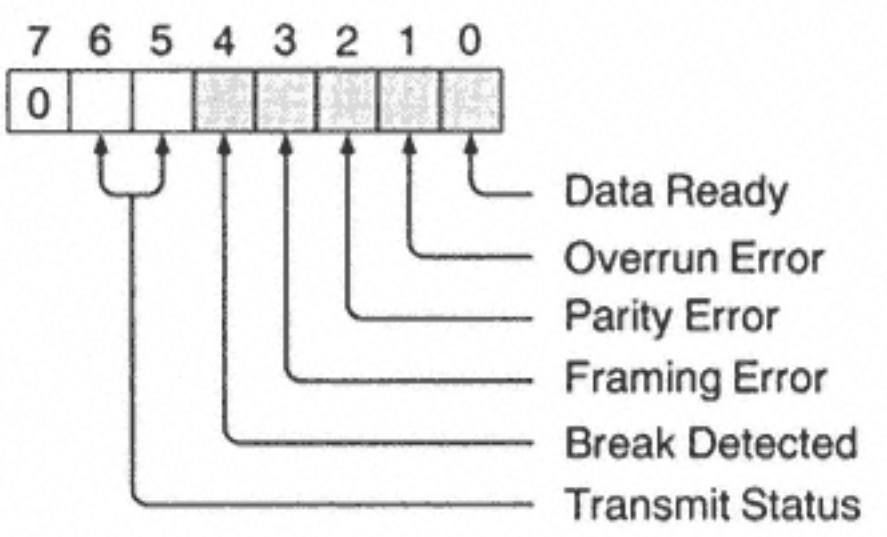

Figure 3.35
Line Status Register: receive status bits.

Data can be read from the UART and displayed using the following polling routine:

```
#define RCVR     0x3F8
#define LSR      0x3FD
#define DA       0x01
main()
{
/* Initialize UART */
for (;;)
{
      while ((inportb(LSR) & DA) == 0);
      printf("%c \n", inportb(RCVR));
}
}
```

Port `0x3F8` supports both transmission and reception because the processor can never perform both operations simultaneously and the UART can determine action being performed upon the port (i.e., read or write). If port `0x3F8` is being written to, it means that a byte is to be transmitted; if the port is being read from, it means that the byte in the receive register is to be supplied to the processor.

Finally, before the UART is used, as part of the initialization procedure both the Line Status Register and the Receive Buffer Register should be read, thereby ensuring that meaningless data or old line-status values are ignored:

```
void clear_regs()
{
   /* Clear registers */
   (void) inportb(LSR);
   (void) inportb(RCVR);
}
```

UART reception handling with interrupts is described later in this chapter.

UART Errors

The Line Status Register indicates more than whether a byte can be transmitted or a byte is available; four other conditions are detectable as well:

Overrun error. The UART reads and places a byte in the Receive Buffer Register before the processor has read the original byte in the Receive

Buffer, thereby overrunning the original byte. The Overrun Error bit is set if an overrun error has occurred; it is cleared by the processor reading the Line Status Register. The original byte is lost.

Parity error. An incoming byte is received with a parity error. The Parity Error bit is set if a parity error has occurred; it is cleared by the processor reading the Line Status Register. The processor can read the byte from the Receive Buffer Register.

Framing error. A stop bit is not found when one is expected, causing a framing error and resulting in the UART setting the Framing Error bit. The Framing Error bit is cleared by the processor reading the Line Status Register.

Break detected. The UART has detected a condition where the line has been in the space condition for at least the time it takes to send a single byte in a complete frame (i.e., start bit, data bits, parity bit, and stop bits). This condition is signaled by the UART setting the Break Detected bit and is cleared by the processor reading the Line Status Register. As mentioned earlier, when a break is detected, many 8250s will set the Framing Error bit as well.

Whenever the Line Status Register is read, the processor obtains the value of the register; however, the register's original contents are lost. This means that if two (or more) bits are set (for example, Data Ready and Parity Error), the register must be read and its value saved so that *all* the status bits can be checked. As an example, the following code fragment checks for any occurrence of the four conditions described above:

```
#define LSR     0x3FD
#define DA      0x01    /* Data Available */
#define OE      0x02    /* Overrun Error */
#define PE      0x04    /* Parity Error */
#define FE      0x08    /* Framing Error */
#define BRKFND  0x10    /* Break Found */

void check_lsr()
{
char lsr_value;
lsr_value = inportb(LSR);
if (lsr_value & DA) data_avail();
if (lsr_value & OE) overrun_error();
if (lsr_value & PE) parity_error();
if (lsr_value & FE) framing_error();
if (lsr_value & BRKFND) break_found();
}
```

Line status changes can also be signaled using interrupts.

UART Interrupts

Thus far, all UART access has been done by polling. From the material covered in Chapter 2, it should be clear that polling wastes processing power because the processor could be performing other tasks. To overcome this problem, the UART has been designed to generate four types of interrupts, three of which are now considered (the fourth is discussed in Chapter 8): line status change, data available, and transmit ready.

Before the UART can signal an interrupt, both the UART and the 8259 Interrupt Controller (see Chapter 2) must be initialized. In the case of the UART, the Interrupt Enable Register (port `0x3F9`) should indicate which interrupts the processor will service (note that bits 4 through 7 are ignored) as shown in Figure 3.36.

To enable an interrupt, the corresponding bit in the Interrupt Enable Register must be set. For example, to enable the Data Available and Receive Line Status interrupts, the Interrupt Enable Register has to be set to `0x05` using `outportb()`:

```
#define IER      0x3F9    /* Interrupt Enable Register */
#define DATA_AV  0x01     /* Data available */
#define TX_HR_MT 0x02     /* Transmit holding register
                             empty */
#define RVC_LS   0x04     /* Receive line status */
#define MDM_CHG  0x08     /* Modem status change */

uart_init()
{
/* Other initialization statements */
outportb(IER, DATA_AV+RCV_LS);
}
```

In addition to setting the Interrupt Enable Register, the 8259 must also be initialized so that interrupts from the UART will be made available to the PC. The UART differs from the other devices studied so far in that the PC can support *multiple* UARTs. Accordingly, for each UART on the PC, there must be a device bit mask available in the 8259 Interrupt Controller. In a typical single-UART PC configuration, bit 4 of the Inter-

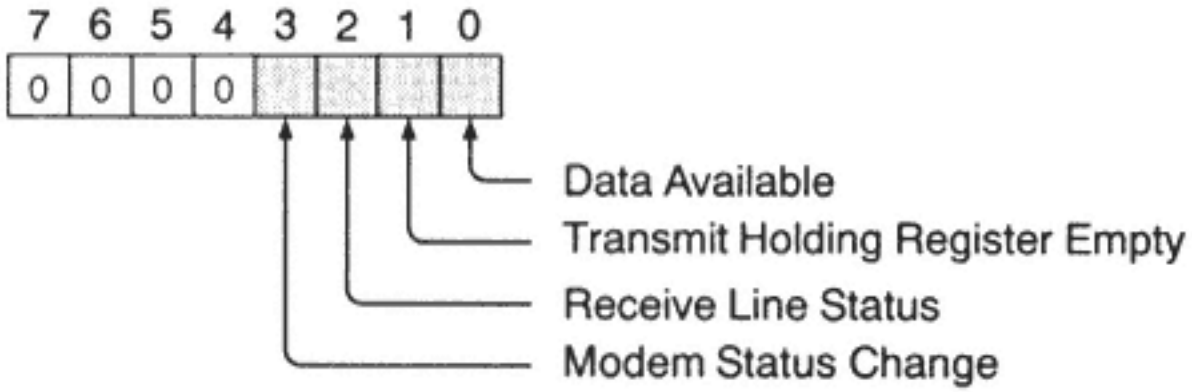

Figure 3.36
Interrupt Enable Register.

rupt Controller is assigned to the UART, while in a two-UART PC, bits 3 and 4 are assigned to the second and first UARTs respectively (see Figure 3.37).

To enable either of these interrupts, the corresponding bit in the Interrupt Controller mask should be cleared. For example, to allow interrupts from both UARTsthe keyboard, and the clockone could write:

```
#define INT_MASK 0x21    /* 8259 Interrupt Control Mask
                            Register */
#define CLKENA   0xFE    /* Clock enable:          11111110  */
#define KEYENA   0xFD    /* Keyboard enable:       11111101  */
#define SP1ENA   0xEF    /* UART 1 enable:         11101111  */
#define SP2ENA   0xF7    /* UART 2 enable:         11110111  */
#define PPENA    0x7F    /* Parallel port enable: 01111111  */
void initialize()
{
/* Other initialization statements */
outportb(INT_MASK, CLKENA & KEYENA & SP1ENA & SP2ENA & PPENA);
}
```

If UART interrupts are *not* required, either the Interrupt Enable Register should be cleared or the UART's bit in the Interrupt Controller mask should be assigned `0xFF`.

The interrupt vectors associated with the two UARTs are 12 (UART 1, addresses `0x30` through `0x33`), and 11 (UART 2, addresses `0x2C` through `0x2F`). `Setvect()` can be used to establish the entry points to the UART interrupt handlers.

Once enabled, the UART can cause interrupts to occur. A UART interrupt only signals the processor that a change has occurred to the UART; it does not indicate the cause of the change. To determine the cause of the interrupt, the processor must read the Interrupt Identification Register, port `0x3FA`, using `inportb()` (see Figure 3.38).

In the Interrupt Identification Register, bit 0 (interrupt pending) is cleared to indicate that at least one interrupt is pending and, if set, signals that there are no other interrupts left to be serviced. Bits 1 and 2 (the interrupt id bits) combine to indicate which of the four possible UART interrupt conditions have occurred. The resulting combinations are assigned priorities, the highest being 6, the lowest 0. The priorities and values of the Interrupt Identification

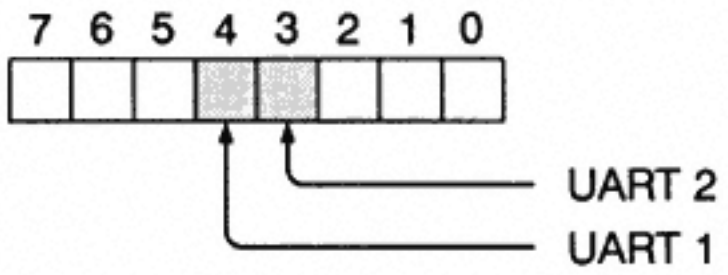

Figure 3.37
UART bit positions in the master 8259.

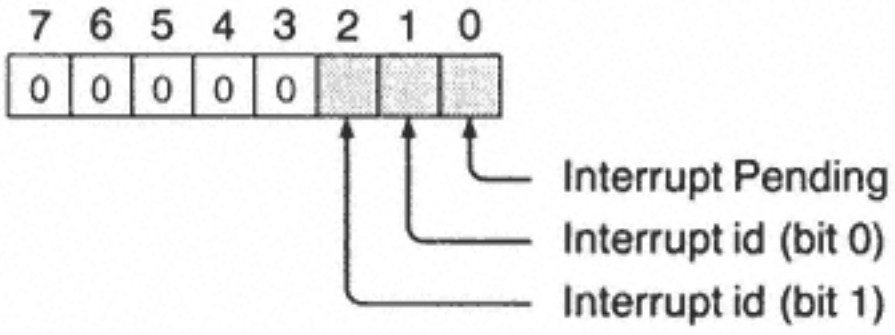

Figure 3.38
The Interrupt Identification Register.

Register are as follows:

Highest (`iir = 6`). A change in the line status (i.e., an overrun, parity, or framing error, or a break interrupt) will cause a Receive Line Status interrupt. The interrupt is cleared by the processor reading the Line Status Register (port S `0x3FD`); the value read indicates the cause(s) of the interrupt.

Second (`iir = 4`). If the UART receives data, the Received Data Available interrupt is signaled. The interrupt is cleared by reading the Receive Buffer Register (port `0x3F8`).

Third (`iir = 2`). The Transmit Holding Register Empty interrupt indicates that the UART is able to transmit another byte. It can be cleared either by the processor reading the Interrupt Identification Register (`0x3FA`) or by the processor writing to the Transmission Register (port `0x3F8`).

Lowest (`iir = 0`). The status of the modem has changed. Modem status change interrupts are discussed in Chapter 8.

Clearing an interrupt condition does *not* necessarily mean that all UART conditions have been serviced: Other interrupts may still be pending. For example, it is possible to clear the line status change interrupt (`iir = 6`), only to be confronted with a received data available (`iir = 4`). Unless *both* of these conditions are cleared during the same interrupt cycle, some UARTs will "lock up" and stop receiving altogether. Therefore, whenever an interrupt is detected and the condition is cleared, the interrupt pending flag of the Interrupt Identification Register must be checked. If it is still 0, other UART conditions are waiting to be serviced.

The 8250 UART is no different from any of the other devices that have been discussed. First, access to the UART takes place through a series of ports; second, the 8259 Interrupt Controller must be signaled that the interrupt is over.

3.5 Point-to-Point Communication Using Commkit

The Commkit module `commkit.c` contains software to control two UARTs: serial port 1 (base port address `0x3F8`) and serial port 2 (base address

0x2F8). The UART base addresses are defined in devices .h; and stored in the array serial_addr (found in commkit.c):

```
#define SP1         0x3F8
#define SP2         0x2F8

int serial_addr[2] = {SP1, SP2};
```

The Commkit software distinguishes between the two UARTs by assigning each a unique queue identifier (SPlIH for serial port 1 and SP2IH for serial port 2). Whenever a UART interrupt occurs, the identifier is used to differentiate between the two serial ports. Since there is no guarantee that the identifiers SP1IH and SP2IH will equal 0 and 1, it is necessary to subtract SP1IH from the supplied serial port number to obtain the correct subscript:

```
int base_addr;
int port_num; /* Set to either SPlIH or SP2IH */
. . .
base_addr = serial_addr[port_num-SP1IH];
```

Once the base address is obtained, all other port register addresses associated with the serial port can be determined.

3.5.1
Initialization

Both UARTs are initialized by the procedure port_init() (found in commkit.c), which takes the identifier (SP1IH or SP2IH) and the line speed of the UART. Port_init() then determines the base address and initializes the following:

The speed of the serial port. Devices.h contains constants for a number of commonly used line speeds and the Divisor Latch Access Bit DLABon:

```
#define DLABon      0x80
#define BAUD96K     0x00c
. . .
#define BAUD50      0x900
```

First, the Line Control Register (base_addr + 3) is set to DLABon. Then, the line speed (baudrate) is written, as a word, to port base_addr.

The line speed is passed to Commkit from the mainline of the program as an integer value. That value can be hardcoded into the program, or the user can be allowed to enter a value from the command line when the program is run. The utility getline_speed() takes a string value and converts it to the equivalent integer line speed value; the integer value is then returned to the calling procedure. The version of get_line_speed()

supplied with Commkit accepts one of six possible values (`50, 300, 1200, 2400, 4800`, or `9600`).

The word length, the number of stop bits, and parity. Constants for setting the Line Control Register are defined in `devices.h`:

```
#define DLABoff     0x00
#define SEVENBIT    0x02
#define EIGHTBIT    0x03
#define WORDLEN     EIGHTBIT
#define ONE_STOP    0x00  /* One stop bit
                             (5, 6, 7, 8) */
#define TWO_STOP    0x04  /* Two stop bits (6, 7, 8);
                             1.5 bits (5) */
#define PEN         0x08  /* Parity enable */
#define EPS         0x10  /* Even parity select */
#define STICK       0x20
#define PARITY      (ONE_STOP+PEN)
#define LCR_DEFN    (DLABoff+WORDLEN+PARITY)
```

The value of `LCR_DEFN` is written to the Line Control Register (`base_addr + 3`).

The enabling of interrupts. The Modem Control Register is initialized using the `devices.h` constants:

```
#define DTR       0x01      /* Data terminal ready */
#define RTS       0x02      /* Request to send */
#define OUT2      0x08      /* Enable serial port
                               interrupts */
#define LOOPBK    0x10
#define MCR_DEFN  (DTR+RTS+OUT2)
```

The constant `MCR_DEFN` is written to `base_addr + 4`. Note that for the UART to generate interrupts, OUT2 must be written to the Modem Control Register as well.

All interrupts are enabled (Data Available, Transmit Holding Register, Line Status, and Modem Status). The constants for each are also defined in `devices.h`:

```
#define DATA_AV     0x01    /* Data available
                               interrupt */
#define TX_HR_MT    0x02    /* Transmit holding
                               register empty */
#define RCV_LS      0x04    /* Receive line status */
#define MDM_CHG     0x08    /* Modem status change */
```

The Interrupt Enable Register (`base_addr + 1`) is initialized to the above constants.

As with all other devices accessed by Commkit, the original interrupt vectors are saved and the new values are written using `setvect()` in the procedure `initialize()`:

```
initialize()
{
   . . .
   old_sp2 = getvect(ll);
   old_spl = getvect(12);
   . . .
   setvect(ll, sp2_ih);
   setvect(12, spl_ih);

   . . .
}
```

`Spl_ih()` and `sp2_ih()` are the entry points for the interrupt handlers associated with each serial port (see below).

The UART Interrupt Handler

Initially, interrupts from either UART cause control to be passed to one of `spl_ih()` (serial port 1 interrupt handler) or `sp2_ih()` (serial port 2 interrupt handler). Both `sp1_ih()` and `sp2_ih()` call a separate procedure, `port_handler()`, which is responsible for processing UART interrupts. The `port_handler()` software distinguishes between the calling procedures by accepting the interrupt handler's identifier (either `SP1IH` or `SP2IH`) as its argument. Figure 3.39 shows how Commkit eliminates the need for two sets of identical software, one for each interrupt handler.

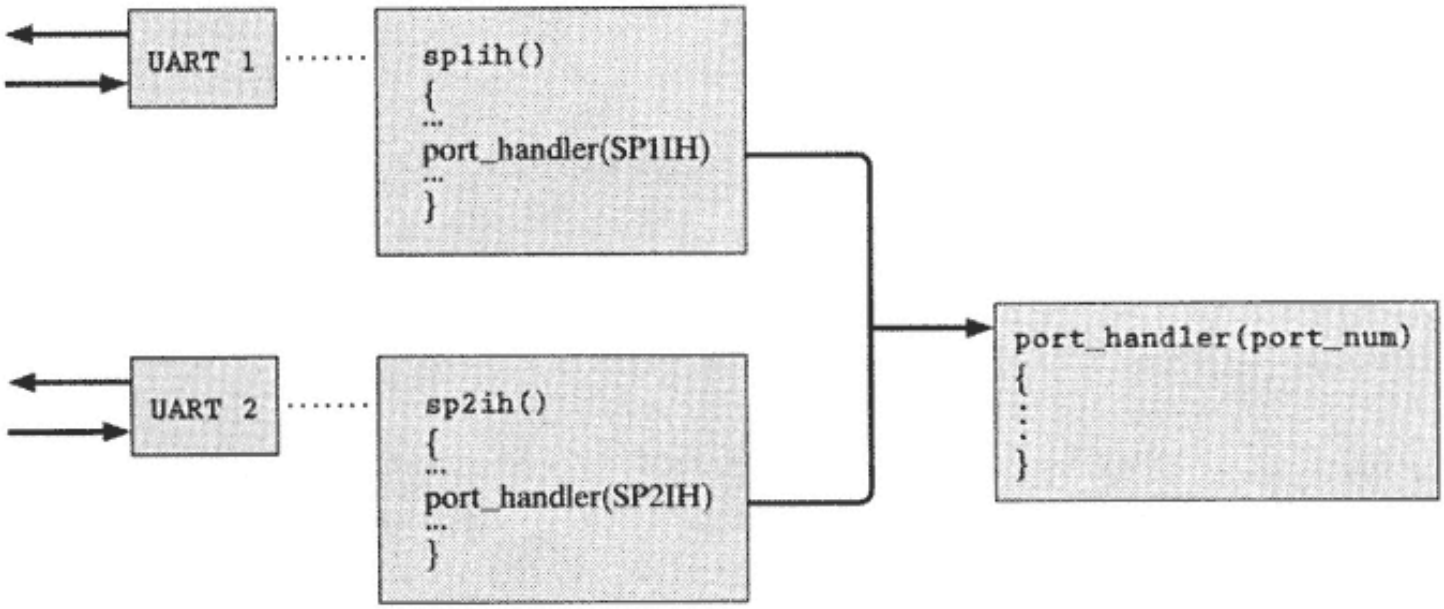

Figure 3.39
How Commkit supports multiple UARTs.

`Port_handler()` obtains the base address of the UART in question from the `serial_addr` array, indexed by `port_num`. With the base address (`address` either `0x3F8` or `0x2F8`), the Interrupt Identification Register (`address + 2`) can be read to determine the type of interrupt that has occurred. By reading the Interrupt Identification Register in a loop, all pending interrupts can be serviced until no more are detected (that is, when `iir = 1`):

```
    while ((iir = inportb(address + 2)) != 1)
    {
    switch (iir)
     {
    case 6: /* Receiver Line Status Interrupt */
            ch = inportb(address + 5) & Ox1E;
            /* Update error statistics */
            break;

    case 4: /* Data Available Interrupt */
            ch = inportb(address);
            low_level(port_num, RECVDONE, ch);
            break;

    case 2: /* Transmit Holding Register Empty Interrupt */
            low_level(port_num, XMITDONE, 0);
            break;

    case 0: /* Modem Status Change Interrupt */
            /* See Chapter on telephones */
            ch = inportb(address + 6);
            low_level(port_num, MODEMSTATUS, ch);
            break;

    default: /* Unknown IIR value */
             /* Update error statistics */
        }
    }
    outportb(INT_CTL, EOI);   /* Reset 8259 for more
                                 interrupts */
```

What `port_handler()` does once the type of interrupt has been determined depends upon the application. For example, whenever a line status change interrupt occurs, `commkit.c` simply maintains a table of error conditions that is updated based upon the error detected. The table, `errors` is a two-dimensional array (one for each port) defined in `commkit.c` and shown in Table 3.7. The contents of `errors` can be displayed using `serial_port_stats()`, specifying the starting line number as the argument.

Table 3.7 Contents of the `errors` array.

Index	Meaning
`overrun.`	An overrun error has occurred
`parity.`	A parity error has occurred
`framing.`	A framing error has occurred
`break_int.`	A line break condition was detected
`modem.`	A modem status change has occurred
`unknown.`	An unknown `iir` value was detected

To accept calls from the serial port software, the procedure `low_level()` must be extended to recognize the following:

`device`. The serial port identifier (one of `SP1IH` or `SP2IH`).
`code`. The event associated with the device; one of
`RECVDONE`: A data available interrupt has occurred.
`XMITDONE`: The transmit holding register is empty.
`MODEMSTATUS`: A modem status change interrupt has occurred.
`MSG_AVAIL`: A message is available for either of the serial ports (call from `send()`).
`data`. The value of the register read after the interrupt has occurred. For `RECVDONE`, data is the incoming character; for `MODEMSTATUS`, it is the value of the Modem Status Register. The value is `0` for `XMITDONE` and `MSG_AVAIL`.

3.6 Example

In the following section, a PC telephone utility is designed and implemented using Commkit.

3.6.1 The Problem

Many multiuser computer systems support "phone" utilities that allow any number of users to carry on a conversation via asynchronous terminals. The problem addressed here is how to support a similar feature that permits electronic communications between two people sitting at a pair of interconnected PCs.

The requirements are as follows:

Whatever is typed on one terminal should also appear on the other (i.e., remote) terminal.

Each screen should be divided in half, with the upper half (10 lines long) reserved for the local user's input and the lower half reserved for the remote user's output.

The two PCs are connected via their serial ports. The configuration is shown in Figure 3.40.

When a character is entered, it should be displayed on the local user's screen as well as transmitted. Characters received from the asynchronous channel should be displayed in the remote user's half of the screen.

The end-of-session is indicated by either user typing TRL-C. This should cause the program to terminate, returning to MS-DOS. If the remote user signals end-of-session, the local user should receive the message `Remote user signed off', at which point the session ends and control should return to MS-DOS.

Once the local user's half of the screen is filled (either with characters or one too many carriage returns), it should be cleared. Similarly, when the remote user's half of the screen is filled (either with characters or one too many carriage returns), it should be cleared. There is no end-to-end signaling (i.e., the DTEs do not explicitly signal one another to clear the screen). It is assumed that both stations are in "lock-step," so that when one station's software clears its local user's half of the screen, it is assumed that the same action is taking place on the other station, with the remote user's half.

3.6.2
The Design

The problem can be broken down into two distinct parts: handling characters from the keyboard and handling characters from the communications channel. To simplify the design, the foreground process will handle the routing (i.e., from keyboard to channel and vice versa) as well as the display of characters.

Input (from Keyboard)

Characters entered by a user from the local keyboard are displayed on the local screen and forwarded to the remote station. Part of the problem has already been solved, since keyboard input is handled by the keyboard interrupt

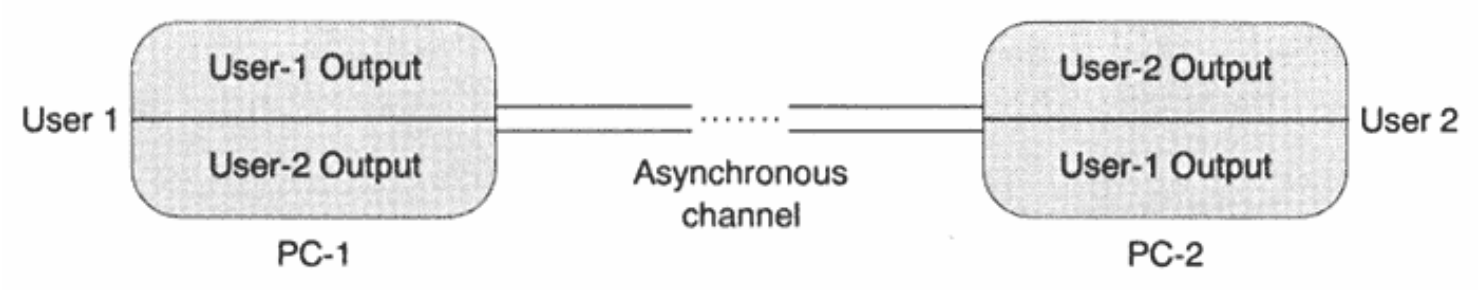

Figure 3.40
Commkit's PC telephone utility configuration.

handler (`keyboard_handler()`) calling `low_level()` with the character. `Low_level()` can then forward the character to the process associated with the `APPLICATION` queue.

Upon receipt of a character from the keyboard interrupt handler, the foreground process examines it. If the character is an *ETX*. it should be forwarded to the remote DTE; once it is forwarded, execution should stop. All other characters must be displayed locally and then forwarded to the remote DTE. However, before the character can be displayed, a check must be made to ensure that there is sufficient space on the screen.

The overall flow of information from the keyboard to the foreground process is represented diagrammatically in Figure 3.41.

The foreground process requires two data structures for screen management: the line and column position of the incoming character. The foreground process algorithm can be written as follows:

1. Accept a character from the `keyboard_handler()`.

2. Check the character. If *ETX*:

    a. Forward *ETX*

    b. Stop the process and return to MS-DOS.

If not *ETX*:

    a. Clear the local user's half of the screen if no positions are available and reset the line/column position.

    b. Display the character in the next line/column position available on the screen.

    c. Forward the character.

Input (from Remote Station)

Characters from the remote station are received by the serial port interrupt handler to which the cable is connected. The interrupt handler forwards each character to the foreground process for display, as illustrated in Figure 3.42.

Since the remote DTE does not indicate when its half of the local screen should be cleared, it is the responsibility of the foreground process to maintain data structures for the line and column position of the incoming remote characters. Initially, the first available character position is in the first column of the lower half of the screen. This position is updated each time a character

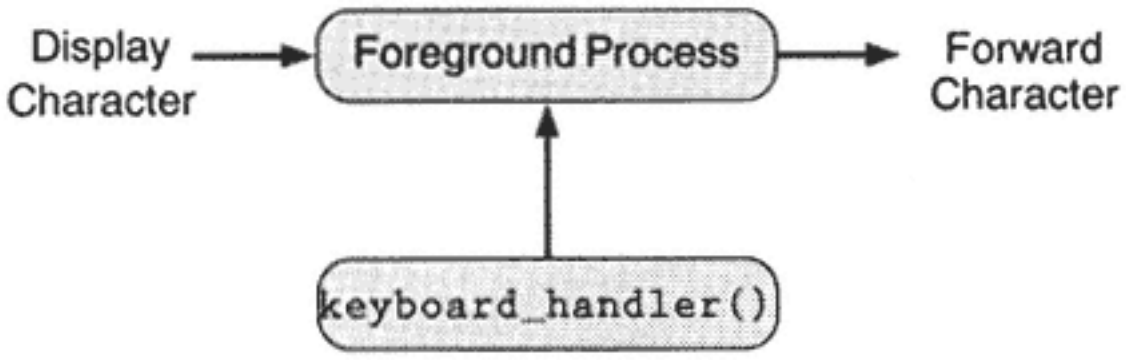

Figure 3.41
Keyboard information flow.

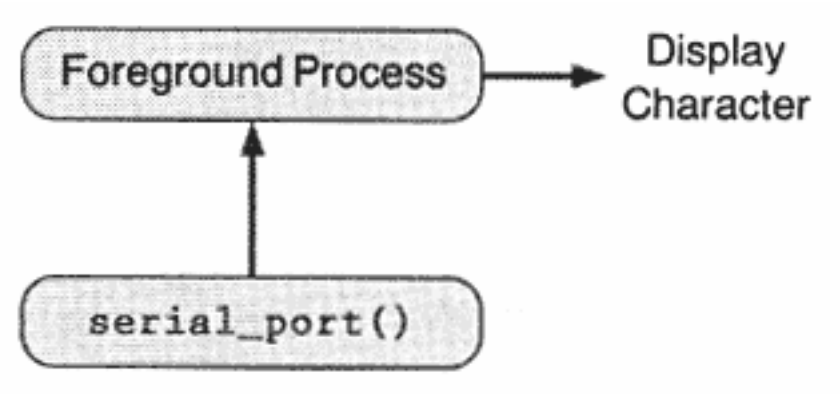

Figure 3.42
Serial port information flow.

is displayed. Using this information, the foreground process can then either display the character or erase the screen.

The algorithm for handling characters from the serial port can be written as follows:

1. Accept the character from `serial_port()`.

2. Check the character. If *ETX*:

    a. Stop the process and return to MS-DOS.

If not *ETX*:

    a. Check whether a space exists on the current screen; if not, clear the remote user's half of the screen and reset the line/column position.

    b. Display the character at the line/column position.

3.6.3 Implementation

An implementation of the point-to-point telephone utility can be found in the file `pt2pt.c` on the Commkit diskette. The software is divided into two distinct parts: the low-level routines for keyboard and serial port handling, and the foreground process for character routing and displaying.

Low-Level

Three procedures support the low-level software. They are described in the following paragraphs.

`system_initialize(). System_initialize()` clears the screen and proceeds to divide it into upper and lower halves to separate local and remote output. The variable `handler_id` is initialized to either `SP1IH` or `SP2IH`, depending upon the serial port through which the communications are to take place.
All software that is called by Commkit must include a routine called `system_initialize()`. `System_initialize()` is called from `initialize()` in Commkit.

`low_level()`. All data from the interrupt handlers and newly arrived messages from the foreground process result in `low_level()` being called. Data from the keyboard and serial port `(RECVDONE)` are enqueued on

the APPLICATION queue; clock interrupts are not sent to the foreground process. When a serial port transmission interrupt has completed, get_char() is called to check the serial port queue for further characters to transmit. Messages from the foreground process are signaled by a code of MSG_AVAIL being received with the serial port's identifier (either SP1IH or SP2IH). If the serial port is idle (indicated by the variable transmitting having a 0 value), get_char() is called; otherwise, get_char() is not called. In either case, transmitting is incremented. Transmitting ensures that messages from the foreground process are not transmitted while the serial port is in the process of sending a previously transmitted character.

get_char(). Get_char() is called when the serial port is no longer transmitting. This is indicated either by an XMITDONE being signaled by the serial port or by a message arriving from the foreground process when the serial port is idle. If a character is on the serial port's queue, it is removed (using low_level_recv()) and transmitted. The variable transmitting is decremented.
To ensure that the telephone utility has time to transmit an *ETX* to the remote station, the global variable running is assigned FALSE after an *ETX* has been transmitted by get_char().

Foreground Process

Three procedures are used to support the foreground process. First, do_pt2pt() waits for characters from either the keyboard or the serial port using the recv() primitive. Characters from the keyboard are passed to do_1cl_scr(), while characters from the serial port are passed to do_rmt_scr(). Control remains in do_pt2pt() until the global variable running is set FALSE either by do_rmt_scr() receiving an *ETX* from the remote PC or by get_char() detecting an *ETX* in the output data stream.

The second procedure, do_1c1_scr(), is responsible for displaying and forwarding locally entered data supplied from do_pt2pt(). It maintains line and column information for displaying each character. Once a character is displayed, the send() primitive is invoked to forward the character to the output serial port.

Finally, do_rmt_scr() displays characters on the lower half of the screen sent from the remote PC. The screen-clearing algorithm in do_rmt_scr() is identical to that found in do_1c1_scr().

Message Flow

The overall message flow in the point-to-point telephone utility is shown in Figure 3.43. The foreground process is shown in two parts simply to distinguish between keyboard input and serial port input. The foreground process

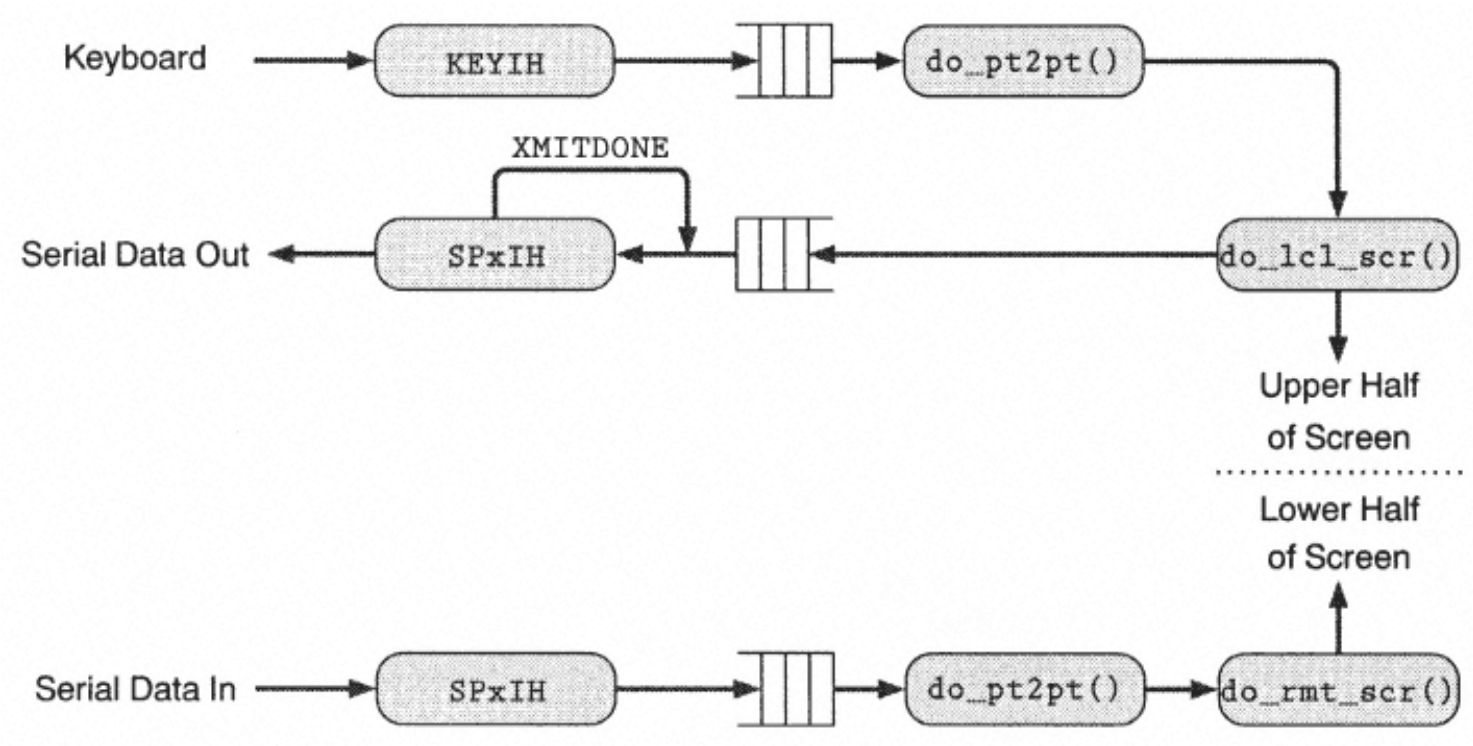

Figure 3.43
Message flow in the PC telephone utility.

waits for messages sent to queue `APPLICATION`: the keyboard and serial port both forward their messages to queue `APPLICATION`.

3.6.4
Using the Telephone Utility

The Commkit diskette contains both source and executable versions of the point-to-point telephone utility. The telephone utility can be run at one of a number of speeds using either of the serial ports.

Hardware Requirements

Two PCs are required, each with at least one serial port. The physical configuration is shown in Figure 3.44. The cable should be connected to the serial ports chosen when starting the program. The local and remote serial ports need not be the same (that is, serial port 1 can connect to serial port 2 or serial port 1, and vice versa).

Creating the Telephone Utility

The telephone utility is created from `pt2pt.c, commkit.obj, srcalls.obj`, and the header files required by Commkit (`general.h, ascii.h`, and `devices.h`). The point-to-point telephone utility `pt2pt.c` can be compiled into an executable using the make utility:

```
C:\> make pt2pt.exe
```

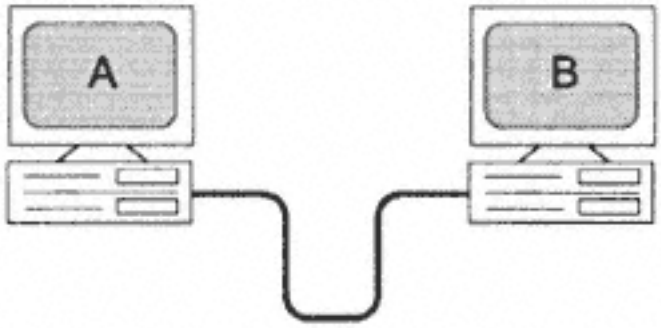

Figure 3.44
Physical configuration of the PC telephone utility.

Running the Telephone Utility

Once compiled, the program can be executed by typing `pt2pt` after the MSDOS prompt, followed by the line speed associated with the port (50, 300, 1200, 2400, 4800, or 9600 bits per second) and the port through which the communication takes place (1 or 2, the default is 1):

```
C:\> pt2pt LineSpeed Port
```

Should the `LineSpeed` be incorrectly entered or the number of words entered on the command line not equal 3, a diagnostic will be generated explaining the required format of the command line. As an example, to set the line speed to 9600 bits per second and to connect the PC through serial port 2, type

```
C:\> pt2pt 9600 2
```

The screen will clear and reappear divided, with the upper half for display of local input and the lower half for display of remote output. Whatever is typed should appear on both the local and the remote PC.

The software forwards characters received from either serial port to queue `APPLICATION`. This can mean that if the cable is connected to a port other than the one selected when running `pt2pt`, characters will still be received and displayed, but characters entered at the local station will not be transmitted. The solution is to connect the cable to the port selected in the command line.

Finally, should the connection be broken for any reason and either of the users continue to enter characters, the local and remote stations will be out of step since the number of characters sent by the PC does not equal the number entered by the user.

3.7 Summary

This chapter examined asynchronous point-to-point communications and how they can be achieved using a device known as a Universal Asynchronous Receiver/Transmitter (UART).

An asynchronous communication is one in which data, in the form of bytes, is sent from one DTE to another at random, unpredictable intervals. All bytes are sent serially (i.e., one bit at a time) between DTEs communicating asynchronously. To distinguish between the communication channel idling and a byte, the bits making up a byte are enclosed in a frame consisting of one start bit and one or more stop bits. Asynchronous communications also offer a limited form of error detection known as parity checking. There are five types of parity: none, mark, space, odd, and even.

The UART handles the conversion of bytes from parallel (within the DTE) to serial for transmission on the communication channel, as well as handling serial-to-parallel conversion for reception. It is also responsible for framing and error detection. To a programmer, a typical UART consists of at least the following:

A transmission register, through which the program supplies bytes for transmission to the UART.

A receive register, from which the program copies the bytes received by the UART.

A status register, indicating the UART's status and including information such as whether

The transmitter is ready.

A byte has been received.

An error has been detected (parity, overrun, or framing).

A break condition has been detected.

A control register, which allows the programmer to specify the number of stop bits, whether parity is required (and what type), the number of bits per character, and the line speed.

For communications to take place between two UARTs, both must agree upon the number of bits per character, the speed of the transmission, the number of stop bits, and the parity. Most UARTs allow these options to be programmed.

3.8
Exercises

1. (A) What would be the value of the missing bit in order to get the specified parity (the missing bit is indicated by a question mark):

1	0	?	0	1	1	0	0	Even
0	0	0	1	0	?	0	1	Mark
1	0	0	1	1	1	1	?	Odd
0	0	0	0	?	0	0	0	Space
1	?	1	1	1	1	1	0	None

2. (A) Does the width of each bit get larger or smaller as transmission speeds increase? Does the time between the transmission of asynchronous

characters increase or decrease as transmission speeds decrease? Explain the result.

3. (A) Describe the two types of overrun that a UART can experience.

4. (A) Calculate the number of data bytes sent in one minute if a UART is to send at 9600 bps with odd parity and two stop bits. Does the number change if even parity is used? Does the number change if one stop bit is sent?

5. (C) All characters except CTRL-C (*ETX*) can be transmitted using the telephone utility; explain why this is so. Propose, implement, and test a method to allow *ETX* to be sent.

6. (C) Extend the sample telephone utility to permit entire lines to be transmitted rather than one byte at a time. Suggest a method to ensure that the line transmitted is received correctly by the remote PC.

7. (C) Modify the telephone utility to emulate a VT-100-type terminal. Use this program to connect to a remote host. It will be necessary to have access to a description of the control characters used by the terminal in order to do this exercise.

8. (C) Determine how sensitive a UART is to minor changes in line speeds. Have one UART transmit a string of characters and a second one receive them. Now, modify the line speed slightly on the receiving UART. At what point does the data become garbled?

9. (A) Show that the recommended divisors for 300 and 1200 bits per second actually produce the required line speeds.

10. (C) Write a small program to demonstrate that the contents of the UART's receive buffer remain unchanged regardless of the number of times the buffer is read.

11. (A) Show, by means of a series of examples, that parity checking (even or odd) will only detect an *odd* number of errors.

12. (A) Explain the conditions whereby mark parity will be equivalent to even (or odd) parity and space parity will be equivalent to even (or odd) parity.

13. (A) In Section 3.3.2, the number of samples on the line performed by the UART each second was calculated for a 9600-bps channel. Calculate line speeds for 2400-, 2000-, and 1200-bps channels; do they all work out exactly using the 1.8432 MHz clock?

14. (A) Parity calculations were performed in this chapter by the exclusive or'ing of the outgoing data bits. Show that the same result can be achieved using exclusive nor'ing.

15. (C) Modify the PC telephone utility to support the following:

> a. CTRL-B (*STX*) indicates that a session is about to begin. As soon as the telephone utility is "activated," it should start transmitting a series of *STX* characters (about once a second). Upon receipt of an *STX*. the software should respond with another *STX*; inform the user that a connection has been made; and then start the session.
>
> b. CTRL-H (*BS*) signals that a character is to be deleted. When a user types *BS*. the character to the left of the cursor should be deleted.

Any character, including carriage returns, can be deleted; this means that a user can delete all of the characters on the local half of the screen. *BS* should be ignored if there are no characters left on the screen.

If a character can be deleted, the *BS* should be transmitted to the remote PC. The remote PC should delete the character to the left of the last character displayed.

c. CTRL-L (*FF*) signifies that the user explicitly wishes to clear the local half of the screen. Upon detection of a *FF*. the upper half of the local user's screen should be cleared regardless of the amount of text on the screen.

FF should also be sent to the remote PC; when the *FF* is received, the remote user's half of the screen should be cleared.

The screens should remain in step after the transmission/reception of either the *FF* or the *BS*.

16. (C) Modify the `commkit.c` routine `get_line_speed()` so that it will take *any* number as a line speed and convert it to the equivalent 16-bit line speed value. Devise a method to test the implementation.

17. (C) Write code fragments to show how the 8250 UART would be initialized for each type of parity (none, even, odd, mark, and space).

18. (C) If an `XMITDONE` interrupt is lost, all further communications will cease since `get_char()` is only called after an `XMITDONE` interrupt. As messages arrive from the foreground process (signaled by `MSG_AVAIL`), `get_char()` will not be called as long as there are messages on the queue (indicated by `transmitting` being greater than 0). Add software to `pt2pt.c` to handle the situation when an `XMITDONE` interrupt is lost. The placing of an upper bound on a piece of software is sometimes called a sanity check and is used to handle those situations when a signal, such as `XMITDONE`, goes missing. This problem will be discussed in more detail in Chapter 5.

4
Physical Connections

4.1
Introduction

Up to this point, any device that can communicate with another device has been called a DTE. Internally, all DTEs are assumed to represent information in the same way: as series of bits with individual values of either 0 or 1. As shown in Figure 4.1, the transmission of information between DTEs is presumed to take place through a UART, across a serial communication channel of an unspecified length, to the other DTE.

Although this is a perfectly reasonable view of how communications take place, little or no consideration has been given to the following issues:

The technology chosen by the manufacturer of a DTE for the storage and movement of information may not be compatible with the technology chosen by another manufacturer. For example, one manufacturer may choose bipolar circuitry (such as TTL, or Transistor-Transistor Logic), while another may choose a metal-oxide semiconductor (such as CMOS, or complementary metal-oxide semiconductor). Connecting a TTL circuit to a CMOS circuit is not impossible; however, there are a number of technical issues that must be resolved, including speed and power (modern CMOS devices operate at higher speeds and consume less power than do TTL

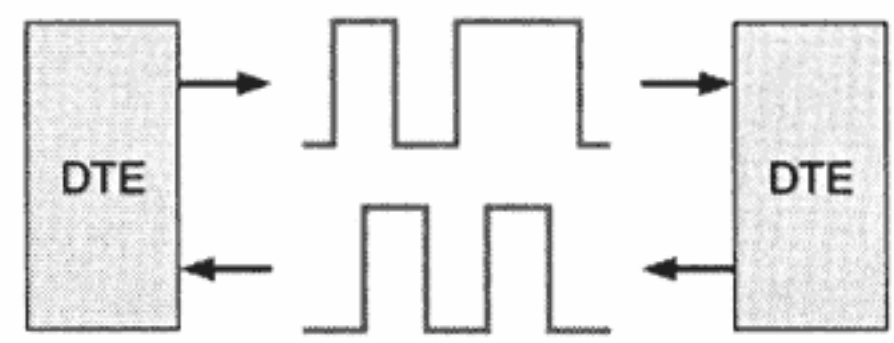

Figure 4.1
The transmission of information between DTEs.

devices) as well as differences in the voltages used to represent the 0s and 1s.

Even if two DTEs are electrically compatible, the electrical signals used within the DTE may have only sufficient power to propagate a limited distance. For example, TTL signals cannot be sent more than about one meter, meaning that the maximum separation between the two DTEs can be no more than that.

This chapter presents a brief introduction to the electrical characteristics and the representation of information in some of the technologies presently available to support the transfer of information between physically connected DTEs.

4.2 Interconnecting DTEs

From the discussion in the previous section and from everyday experiences with DTEs, one can conclude that there are three issues that must be addressed to ensure that information can be transmitted between the interconnected DTEs:

The information within the transmitting DTE must be converted to a signal that can reach the receiving DTE.

The signal must be carried across a channel (a medium such as a wire, optical fiber, or air) to the receiving DTE.

Before supplying the signal to the DTE, the signal must be converted to the internal format used by the receiving DTE.

The first and last of these issues (signal conversion) are performed as part of the DTE in a device known as a level converter, which encodes the information into an agreed-upon signal, transmitted across a medium, that the receiving level converter decodes into the equivalent value on the receiving DTE.

Level converter, information encoding, and transmission media are examined in this section.

4.2.1 Level Converters

At large musical events, singers' voices aren't loud enough to be heard by everyone; instead, they use microphones to amplify their voices to a level audible to the crowd.

Similarly, since the binary information within a DTE typically cannot propagate over distances of more than about a meter, a device that changes the information into a set of signals that can reach and be recognized by the receiving DTE is necessary. To avoid having unique conversion devices for

all possible DTEs, the information is converted into signals agreed to by the various manufacturers. Once the signal reaches the remote DTE, it can be converted to the internal binary representation used by the receiving DTE.

Figure 4.2 shows how the conversion of the information to the signal required on the channel takes place between the UART and the physical connection, with the remote DTE using a level converter.

The level converter is responsible for changing the DTE's internal binary information to that agreed to by the various DTE manufacturers for use on the channel. For example, a mark (1) may be converted to 10 volts, while a space (0) may be converted to +10 volts.

Upon reception of a signal from the physical connection, the level converter must change the signal into an internal value acceptable to the DTE. Continuing with the example, a physical signal of-10 volts would be changed into a binary 1 and a signal of +10 volts would be changed into a binary 0.

4.2.2
Information Encoding

When converting the information within the DTE to a signal that can be transmitted to another DTE, several factors must be considered. First, by increasing the line speed, more information can be transmitted. Second, an increase in the line speed increases the probability that a bit will be received in error. Fortunately, there are a number of techniques that allow the line speed to be increased while reducing the chances of receiving the signal in error.

Non-Return to Zero

The simplest type of encoding is known as Non-Return to Zero-Level (NRZ-L), in which each bit value is associated with a *constant* signal level,

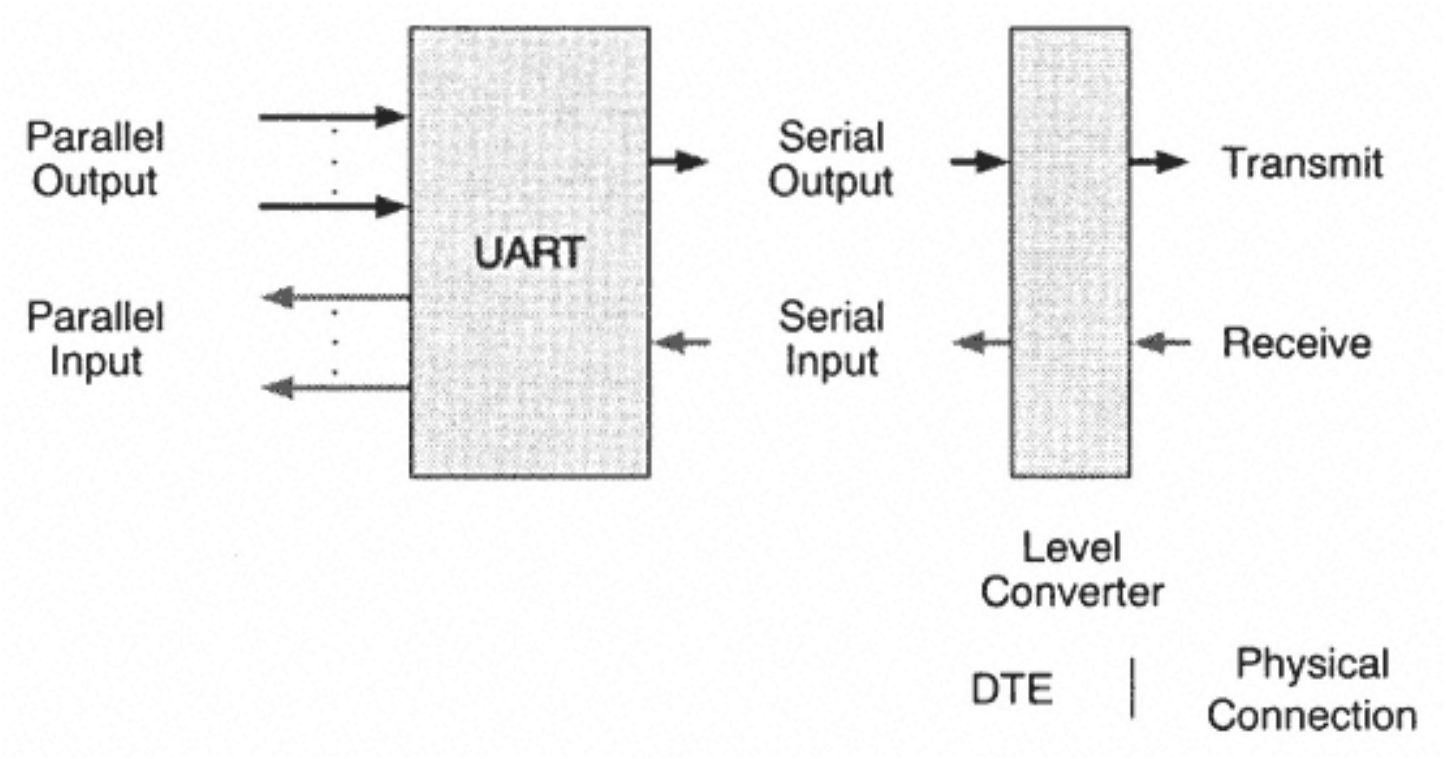

Figure 4.2
The level converter.

either positive or negative (hence the name "NRZ-L," since during the transmission of a bit the signal level remains constant and does not return to 0). For example, a mark may be associated with a negative voltage and a space with a positive voltage. The only transition that occurs is when the signal changes from a mark to a space (or vice versa).

NRZ-L can be implemented with little difficulty. For example, the character *S* transmitted using NRZ-L would appear as shown in Figure 4.3. Remember, the rightmost bit is being transmitted first in this and subsequent examples; all of the examples in Chapter 3 used NRZ-L.

Variations on NRZ include NRZ-I (or Non-Return to Zero-Inverted), in which a *transition* from one level to another indicates a mark rather than an absolute value of signal, as in NRZ-L. The lack of a transition indicates a space; the presence or absence of a transition can be detected by sampling the middle of each bit and comparing it with the level of the previous bit. Figure 4.4 shows how the byte *S* would be transmitted using NRZ-I, in which a transition occurs each time a mark is transmitted. As the line is sampled at the receiving DTE, a transition is taken to indicate a mark; the absence of a transition indicates a space.

NRZ-I is also known as NRZ-M (Non-Return to Zero-Mark). In NRZ-S (Non-Return to Zero-Space), a transition indicates a space, the lack of a transition a mark.

Since NRZ-I and NRZ-S do not refer to an absolute signal value but rather to the presence or absence of a transition, they can offer better reliability over noisy channels than NRZ-L.

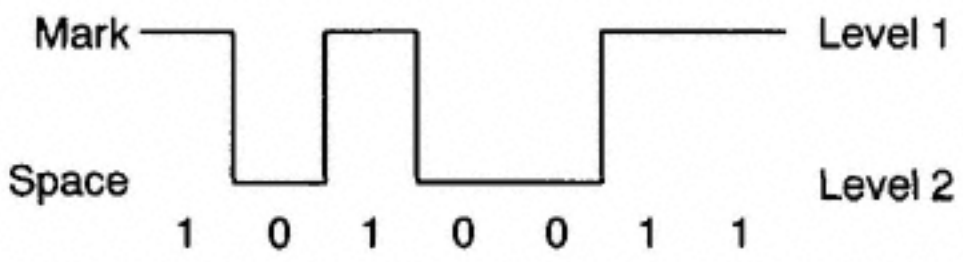

Figure 4.3
Non-Rreturn to Zero-Level (NRZ-L).

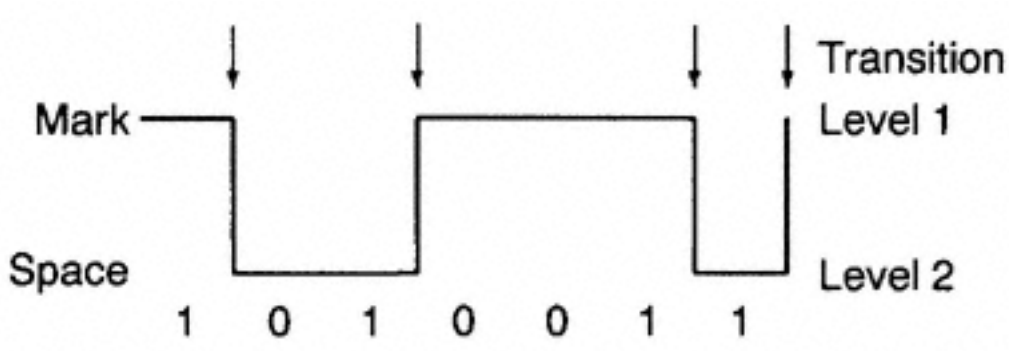

Figure 4.4
Non-Return to Zero-Inverted (NRZ-I).

Return to Zero

A major limitation in the NRZ signal-encoding technique is that if there is a long sequence of 0s or 1s (in NRZ-L), 1s (in NRZ-S), or 0s (in NRZ-I), a constant signal level is produced. A constant signal level can result in the transmitting and receiving clocks losing synchronization.

This problem can be minimized by putting the transition in the *middle* of the bit and requiring the UART to maintain its bit synchronization from the value of the transition (i.e., from high to low or from low to high). A simple example of this encoding technique is Return to Zero (RZ).

Return to Zero uses two signal levels: 0 and non-zero. A space can be represented by a 0 value (i.e., no transition), whereas a mark requires *two* transitions in the time it takes to send one bit. The first transition (from the 0 signal level to the non-zero) occurs at the *start* of the mark bit, while the second (from the non-zero signal level to the 0) occurs in the *middle* of the mark bit. The last half of the mark bit is the 0 signal level. Return to Zero ensures that the receiving DTE receives a transition (at least for marks), thereby allowing a degree of clock synchronization.

Figure 4.5 shows how the byte *S* could be transmitting using Return to Zero. Although each mark causes a transition, potentially allowing the receiving clock to synchronize with the transmitting clock, a sequence of spaces could still result in the clocks drifting. Note also that in the NRZ-L encoding scheme, there is a one-to-one relationship between the bit and the transition, but in RZ one bit is represented by *two* transitions.

Biphase

Return to Zero attempts, but fails, to reduce the problem of the two stations losing bit synchronization. By extending the Return to Zero encoding technique to include a transition on *every* bit, it is possible to embed the transmitter's clock directly into the data stream. For example, a mark could be represented as a high-to-low signal transition, while a space could be a low-to-high signal transition. This technique, known as biphase encoding, allows the receiving DTE to extract both clocking information (thereby staying in synchronization with the transmitter) and data from the signal received.

As an example, to transmit the byte *S* using a biphase encoding technique in which the first half of the signal is equal to the bit and the second half

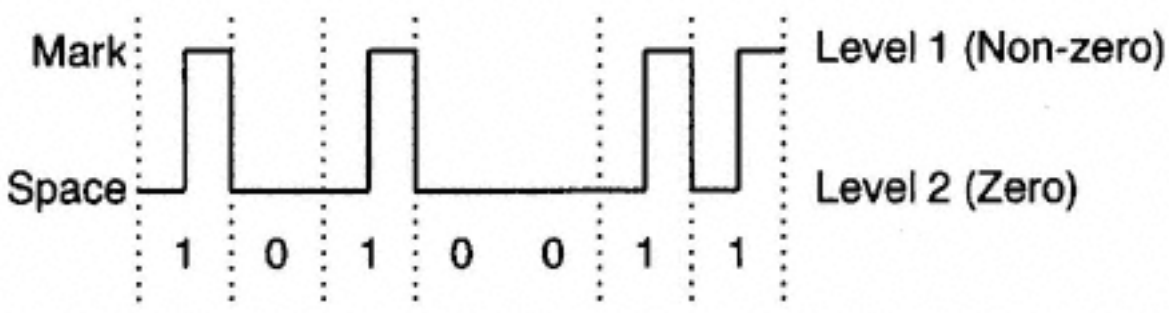

Figure 4.5
Return to Zero.

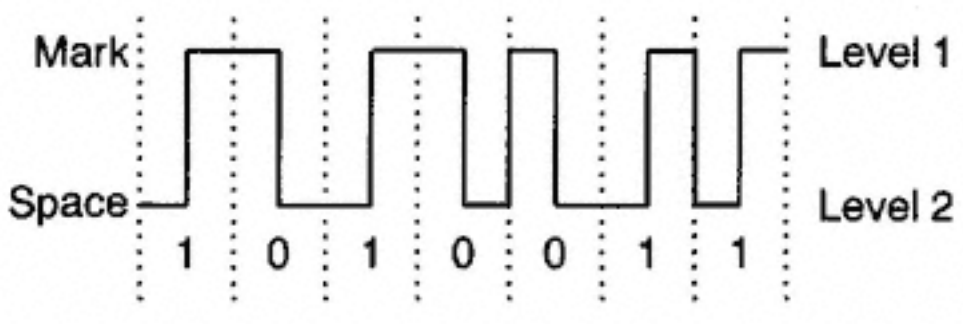

Figure 4.6
Biphase.

is equal to the inverse of the bit, the situation shown in Figure 4.6 would occur. Now every bit causes a transition: a transition from Level 1 to Level 2 indicates a 1, and a transition from Level 2 to Level 1 indicates a 0.

Clock resynchronization can be achieved by the receiving DTE sampling the channel when the center of a bit (i.e., a transition) is expected. If a transition is detected, the clocks are in synchronization; should the transition occur after it was expected, the receiver's clock is running slower than that of the transmitter and must be adjusted. However, if the transition is detected before the receiving DTE expected it, the receiver's clock is running faster and must be adjusted accordingly.

Probably the best known example of biphase encoding is Manchester encoding, which is widely used in high-speed communications. Manchester encoding will be discussed in greater detail in subsequent chapters.

4.2.3
Transmission Media

Once the signal has been generated by the level converter, it is transmitted through the communication channel to the remote DTE. The communication channel consists of a medium that is either bounded (such as wires or optical fibers) or unbounded (the air or the vacuum of space). Some of the more common media are now considered.

Two-wire Open Lines

A two-wire open line consists of two wires insulated from each other, one carrying the signal (typically a voltage or a current), and the other carrying the ground reference. The receiving DTE determines the value of the signal by comparing the signal with the ground reference.

Two-wire open line systems are best used in applications requiring low bit rates (less than 19.2 kilobits per second) and limited distances (less than 50 meters). They are susceptible to electromagnetic interference: should the noise affect only one of the two wires, the receiving DTE can misinterpret the signal. Consider, for example, the situation in which the ground reference is taken as 0 volts, with mark being represented by a voltage greater than the ground reference, and space by a voltage that is less than the ground reference. If a mark is transmitted and noise causes the ground reference to

be greater than the voltage used for mark, the receiver, taking the difference, could interpret the mark as a space (see Figure 4.7).

Twisted-Pair

Considerable improvements can be made upon two-wire systems by taking the pair of insulated wires and *twisting* them together. Any electromagnetic interference should affect *both* wires (not just one). If the receiving DTE determines the value of the signal from the difference between the signal and the ground reference, and if both wires have been affected by noise, then, ideally, the difference will be the same.

The resulting benefits of twisted-pair over two-wire include potentially higher line speeds and longer transmission distances (line speeds in the range of 1 million bits per second at distances of up to one kilometer can be achieved). Since twisted pairs are widely used by telephone companies for voice communications, they are often bundled in groups of up to 3000 pairs in a single cable. To reduce the likelihood of interference (or crosstalk) between adjacent pairs, the twist length is varied within the cable. However, when bundled, twisted pairs are still subject to interference; as the frequency increases within a twisted pair, the amount of crosstalk on adjacent pairs also increases.

Twisted-pair is another example of a bounded medium.

Coaxial Cable

Although twisted-pair offers numerous advantages over two-wire open line systems, its limited bandwidth results in signal attenuation at higher frequencies. This limitation can be overcome using coaxial cable, another bounded medium.

Coaxial cable consists of a signal conductor and a ground reference conductor separated by a dielectric material (a material that does *not* conduct electricity) and encased in an insulating material. One of the conductors is a wire that passes through the center of the dielectric, and the other is wrapped around the dielectric (typically as a wire braid). Signals transmitted by coaxial cable are less prone to the effects of electromagnetic radiation and signal attenuation. Consequently, much higher line speeds are possible, up to 800 million bits per second (Mbps) over distances of about 1.5 kilometers.

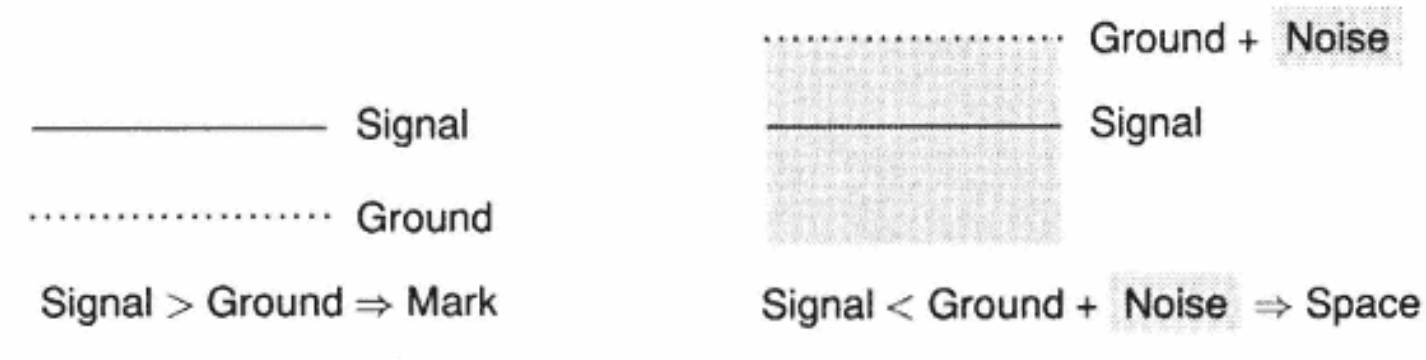

Figure 4.7
The effect of noise on a two-wire open line.

Optical Fiber

All of the bounded media examined thus far are constructed from metallic material and as such are subject to the effects of electromagnetic interference, crosstalk, and limited bandwidths. The development of optical fibers that carry signals in the form of light rather than currents or voltages can overcome many of the limitations associated with metallic media at extremely high bandwidths.

Optical fibers are known as dielectrics and are made from plastic or glass. The optical fiber (the core) is enclosed in a cladding and a sheath. The signals through it are sent as light, generated by light-emitting diodes (LEDs) or injection laser diodes (ILDs). Although ILDs consume less power than LEDs and produce better signals, the cost and lower reliability of ILDs make LEDs more attractive. Signals are detected by the use of a photo-detector.

Finally, it is worth noting that signals in an optical fiber are subject to various forms of attenuation: scattering (the radiation of the signal); absorption (the conversion of light energy to heat); connection losses at joints and splices; and losses at bends in the fiber. However, if treated carefully, optical fiber can operate at speeds of up to 8 *billion* bits per second at distances of almost 70 kilometers.

Microwave

Probably the most common example of unbounded communications is microwave, which is used for the transmission of television, voice, and data signals. Microwave has the advantage over the other techniques previously described in that there is no physical connection between the transmitting and receiving DTEs. Its limiting factor is that it cannot pass through objects, meaning that microwave transmitters and receivers must have *line-of-sight* transmission. It is not uncommon to find microwave in use between buildings.

Communication satellites can also be used for microwave transmission. Ground stations communicate with a geostationary satellitethat is, a satellite that remains in a fixed position above the Earth; as with microwave, the groundstations must be able to "see" the satellite. Transmissions to the satellite are sent on one frequency (the uplink) and returned on a second frequency (the downlink); different groundstations are assigned different uplink and downlink frequencies. The signals received and transmitted by the satellite are *repeated* using a device known as a transponder.

4.3
Standards

Before a communication can take place, both DTEs must agree to a common protocol that covers things such as line speed, word size, parity, and

error-recovery methods. Similarly, there must be agreement as to the physical connection of the DTEs, covering aspects such as the voltages used to represent mark and space, the type of information encoding on the channel, and the physical properties of the channel itself. Protocols used to describe the physical connections between the DTEs are often called standards.

Standards are rules created by groups or organizations in an effort to ensure that equipment from different manufacturers can exchange information. The standards describing the physical connection cover a wide range of issues, including how signals (marks and spaces) are to be encoded and the maximum allowable distance between DTEs.

There are many organizations involved in defining communication standards; however, some of the more widely known for standardizing the physical connections between interconnected DTEs follow:

The Electrical Industry Association (EIA), a U.S.-based organization to which many North American manufacturers of electrical equipment belong. EIA standards are often prefixed by the letters RS, which mean "recommended standard."

The International Telecommunications Union, or ITU (which has superseded the Comité Consultatif International Téléphonique et Télégraphique, or CCITT). This is an organization sponsored by the United Nations that is responsible for producing international telecommunication standards. ITU membership is made up of representatives from national PTTs (Post, Telephone and Telegraph) and governmental organizations. All countries and many companies belong to the ITU.

ITU defines a number of standards for communication systems, including analog (designated by the prefix "V."), digital (designated by the prefix "X."), and ISDN (designated by the prefix "I."). ITU documents are *recommendations*; those intended for telecommunications are denoted ITU-T.

Other standards organizations that will be discussed in subsequent chapters include the International Organization for Standardization (ISO) and the Institute of Electrical and Electronic Engineers (IEEE).

A number of DTE interconnection standards are available, defined by both EIA and ITU. The remainder of this chapter considers one of the most widespread interconnection standards, RS-232-C (and its ITU equivalent, V.24), and briefly examines some other interconnection standards.

4.4 RS-232-C (V.24)

RS-232-C (ITU-T V.24) defines the electrical, physical, and mechanical properties of a widely used (and often abused) EIA standard. Almost all existing DTEs support RS-232-C, including the serial communication ports supplied with most PCs.

4.4.1
Physical Characteristics

An RS-232-C connector, shown in Figure 4.8, has 25 pins and is known as a D-connector because of its shape (one side is slightly longer than the other, making it look like a *D*).

Since not all 25 pins are necessary for a DTE-DTE connection, 9-pin D-connectors are also available (see Chapter 8). Of the 25 pins, only three are of interest at the moment (the remaining pins are discussed in detail in Chapter 8):

Pin 2. Serial data is transmitted on this pin (on a 9-pin connector, the transmit pin is pin 2).

Pin 3. Serial data is received on this pin (on a 9-pin connector, the receive pin is pin 3).

Pin 7. Used to establish signal ground (on a 9-pin connector, signal ground is pin 5).

There are two "genders" of RS-232-C connector: male (the one with pins), and female (the one with holes). In Figure 4.8, the D-connector has female numbering (the male is reversed, so that when the connectors are coupled the numbering corresponds correctly). The RS-232-C standard expects all DTEs to have male connectors, meaning that the cable connecting two DTEs should have female connectors. However, since RS-232-C is a (very) loose standard, it is possible to find DTEs with female connectors. When the cable's connector is the same gender as that of the DTE, a device known as a gender mender is required, which is simply a very short RS-232-C cable with the same gender of connector at either end.

The maximum recommended separation between two DTEs connected by an RS-232-C cable is 50 feet (15 meters), and the maximum recommended bit rate is 9600 bps. However, as with most aspects of RS-232-C, these maximums are often exceeded; for example, it is possible to run 9600 bps at distances up to 250 feet (75 meters). An RS-232-C cable connects each pin at one end of the cable to the same pin at the other end of the cable (i.e., pin *P* connects to pin *P*).

4.4.2
Interconnecting DTEs

Before a pair of DTEs can communicate, they must be connected. In the case of RS-232-C, the standard RS-232-C cable described above will *not* work. The

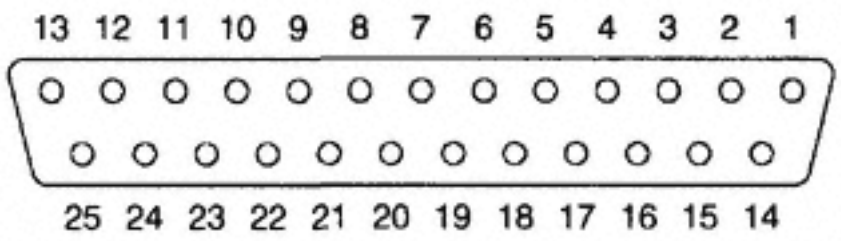

Figure 4.8
The 25-pin RS-232-C (ITU-T V.24) D-connector.

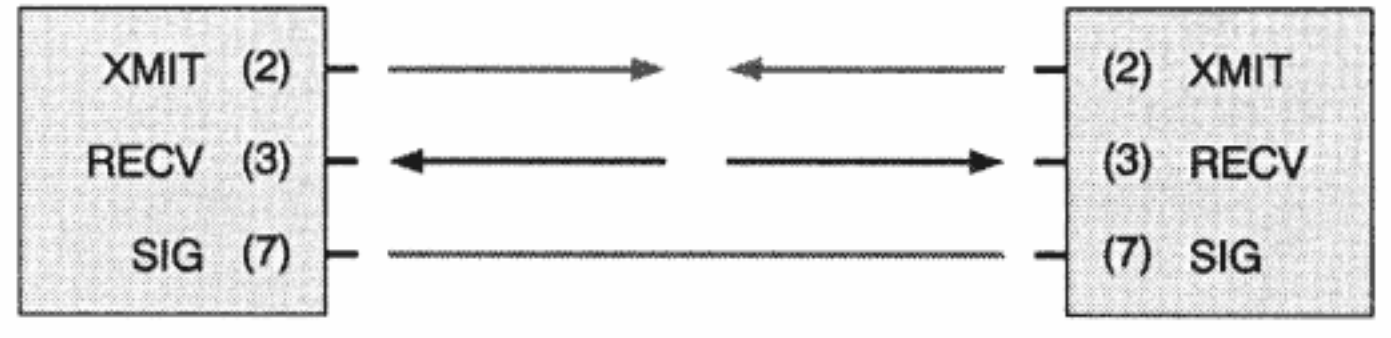

Figure 4.9
The result of directly connecting a pair of DTEs.

reason for this is fairly straightforward: as shown in Figure 4.9, if a DTE is connected to another DTE, both transmit pins will connect through to one another and both receive pins will do the same (assume a 25-pin connector).

To allow one DTE to connect to another (without modifying the pin connections on the RS-232-C port), something must be done to the cable to ensure that the communications can take place. The solution is to *swap* the transmit and receive lines, so that whatever is sent on a DTE's pin 2 appears on the other DTE's pin 3. A cable wired in this fashion is known as a nullmodem cable. The minimum null-modem wiring is shown in Figure 4.10.

4.4.3
Signal Levels

The typical voltage range for RS-232-C is ±15 volts, with a positive voltage (between +3 and +15 volts) indicating that the line is in a space condition and a negative voltage (between 3 and 15 volts) indicating the mark condition (see Figure 4.11). RS-232-C uses NRZ-L for its signal encoding.

4.5
Other Interconnection Standards

RS-232-C is by no means the only or the best interconnection method available. This section examines two other methods. The first, the 20-milliampere (mA) current loop, is quite old and not associated with a particular standard; the second, RS-449, is the proposed successor to RS-232-C.

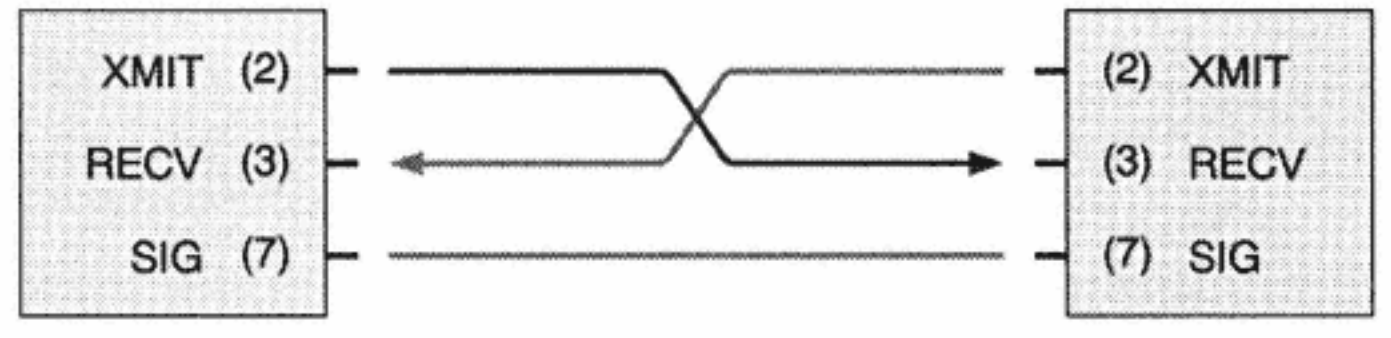

Figure 4.10
The minimum wiring for a null modem.

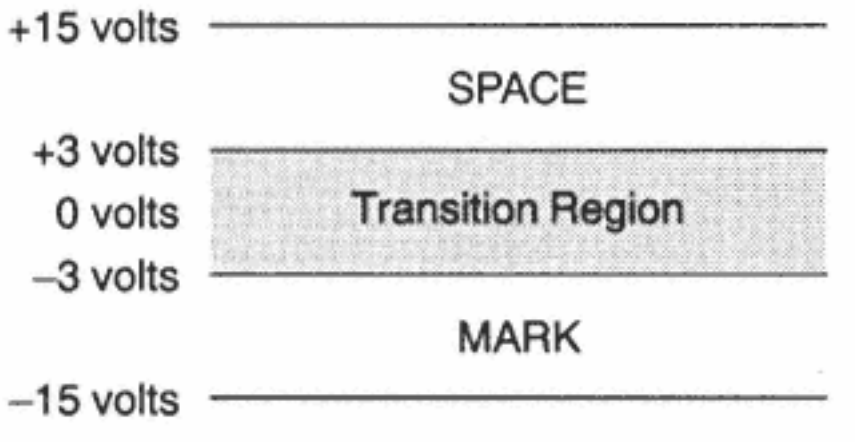

Figure 4.11
RS-232-C signal levels.

4.5.1
20-Milliampere Current Loop

The 20-milliampere current loop predates the computer, since the current loop was (and still is) used in teletype communications. The basic operation consists of sending a 20-milliampere current pulse from the local DTE to the remote DTE: the presence of the current indicates a mark, and the absence indicates a space. The presence and absence of a current to represent mark and space are known as neutral working (a variation, in which the current flows in one direction for mark and in the other direction for space, is known as polar working).

The fundamental components of a 20-mA current loop are a current source, a current switch, and a current detector. Each DTE interface consists of two parts: the transmitter, containing the current switch; and the receiver, containing the current detector. As shown in Figure 4.12, for a full-duplex communication two current loops are required (a total of 4 wires), one for each direction of the data.

The interface that supports the current source is said to be the active interface, while the interface without the current source is the passive interface. Typically, one DTE is active and the other is passive (as in Figure 4.12). Regardless of the configuration, an active transmitter must send to a passive receiver or a passive transmitter to an active receiver. In those situations where both the transmitter and the receiver are "the same," an intermediate device is required either to supply a current source (if both are passive) or to act as a passive receiver and passive transmitter (if both are active).

The 20-milliampere loop is suitable for distances up to about 1500 feet (roughly 450 meters). Crosstalk is one of the disadvantages of using the 20mA loop over long distances.

The main problem with the 20-mA loop is that no standard is defined. Consider two 20-mA DTEs, one designed as an active interface using 400 volts and a 20,000-ohm resistor, and the other designed as a passive interface for 4 volts and 200 ohms. Although both support 20 mA, connecting the two could result in damage to the passive DTE. Figure 4.13 shows how fiber optics can be used to isolate the two DTEs, thereby avoiding the voltage problem.

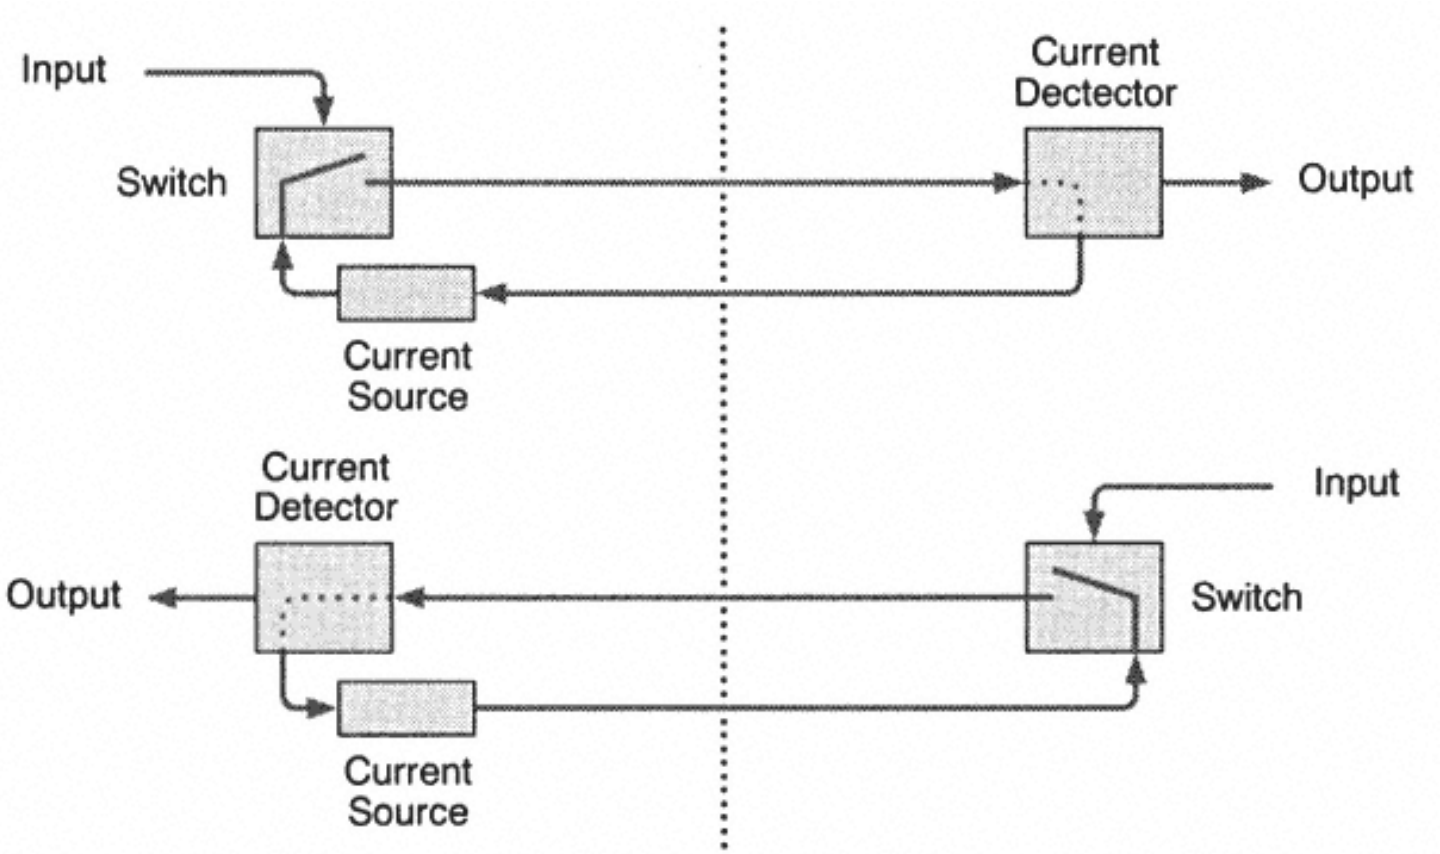

Figure 4.12
A full-duplex 20-mA current loop implementation.

Figure 4.13
Isolating DTEs with fiber optics.

Some PCs support a current loop interface. For example, some asynchronous adapters for the IBM PC have an internal switch that can select either RS-232-C or 20-mA current loop (pins 9 and 11 for the transmit current loop, and pins 18 and 25 for the receive current loop).

4.5.2
Other EIA Standards

RS-232-C, for all its seeming popularity, has a number of serious limitations that restrict the distances over which it can be used. For example, because of line capacitance, it is possible for signals to become distorted to the point where the data is received incorrectly. Similarly, should the ground reference (pin 7) be different for the two DTEs, undesirable electrical characteristics can be applied to the transmitted signal.

In recognition of these and other limitations of RS-232-C, the EIA has proposed three new standards: RS-449, RS-422-A, and RS-423-A. RS-422A and RS-423-A define the electrical circuits that use the physical interface defined in RS-449.

The RS-449 standard defines the physical connections required to support either RS-422-A or RS-423-A. It is procedurally similar to RS-232-C and is

intended to offer an orderly means of migrating from the old standard to the new standard. Some of the differences between RS-449 and RS-232-C are given in the following paragraphs.

The RS-449 standard calls for *two* sets of connectors: a 37-pin connector for data, control, timing, and diagnostics, and a 9-pin connector for a secondary channel circuit. RS-232-C has a single 25-pin connector carrying all signals.

RS-449 supports both balanced and unbalanced circuits; RS-232-C supports only unbalanced circuits. A balanced circuit is one in which the signals are carried between the DTEs on a *pair* of wires. They are sent as a current down one wire and return on the other; the two wires create a complete circuit. An unbalanced circuit is one in which the signal is sent over a single wire, with the DTEs sharing a common ground. Electrical problems can arise if the ground potential differs between the two DTEs.

A balanced circuit is less affected by noise and produces less noise than does an unbalanced circuit. In a balanced circuit, the receiving DTE determines the value of the signal (mark or space) by comparing the difference between the transmitted signal and the ground. Should the line be affected by noise, the difference between the transmitted signal and ground would be the same as if the noise were not present (Figure 4.14). The equivalent ITU-T standard for RS-449 is V.35.

The standards RS-422-A (ITU-T V.11) and RS-423-A are intended for balanced and unbalanced circuits, respectively. Since RS-422-A supports balanced circuits, two separate wires are used for each signal (transmit and receive), doubling the number of circuits but permitting higher data rates (see below). In RS-422-A the differences between the voltages on the two wires determine whether a mark or a space is sent. If the signal difference is positive and more than +0.2 volts, a mark is received, whereas a negative difference of more than 0.2 volts indicates a space. The smaller transition region (0.4 volts in RS-422-A compared to 6 volts in RS-232-C) is through the use of the balanced circuit.

RS-423-A supports unbalanced transmission and is designed, in part, as a way for RS-232-C users to migrate to RS-449, since RS-423-A can support

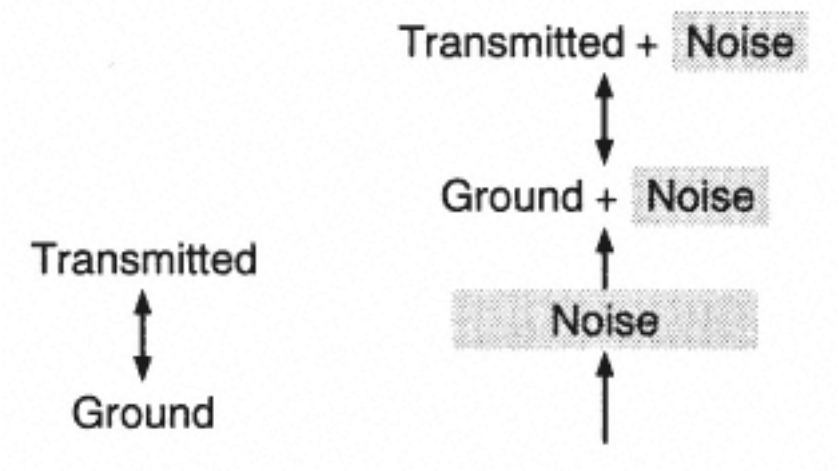

Figure 4.14
The effect of noise on a twisted-pair (RS-449) cable.

both RS-422-A and RS-232-C. In RS-423-A, a mark is indicated by a4-volt difference between the signal and the *common* ground wires, while a space is indicated by a +4-volt difference. These voltages are compatible with the existing RS-232-C standard, thereby permitting RS-423-A interconnection with RS-232-C. (Since RS-423-A receivers also handle the RS-422-A 0.4-volt transition region, RS-422-A transmitters can be used with RS-423-A receivers.)

Table 4.1 highlights the differences between balanced and unbalanced circuits. These are also illustrated when comparing the speeds and distances attainable with RS-422-A and RS-423-A.

4.6 Summary

This chapter introduced some of the concepts surrounding the generation of signals and the technology available for the transmission of data between DTEs. It examined three aspects of the physical connection:

The encoding of signals. Data within a DTE cannot simply be transmitted to a remote DTE, for a number of reasons: The signals might not be strong enough to reach the destination, or the electrical representation of the signals may differ from DTE to DTE. To overcome these problems, devices known as level converters are used to "boost" the signal strength as well as to convert the signal to an encoding agreed upon by both DTEs.

Connecting DTEs. For information to reach one DTE from another, some form of connection is required. In all of the applications considered thus far, the connection has been physical (typically copper wires); however, there are alternatives to wire, including optical fibers and microwave.

Standards. Standards are needed to ensure that DTEs from different manufacturers can be connected. These standards define how signals are encoded, the physical wiring of the devices, and even the shape of the connectors. Without these agreements, data communications would be extremely difficult.

In subsequent chapters, RS-232-C will be examined in more detail and additional standards will be discussed.

Table 4.1 Effect of distance on line speed.

Standard	At 1000 Meters	At 10 Meters
RS-422-A	100,000 bps	10,000,000 bps
RS-423-A	3,000 bps	300,000 bps

4.7
Exercises

1. (A) Draw the bit pattern for the transmission of the letter *S* using NRZ-S (space) encoding. Compare this to NRZ-I.

2. (A) Show, by means of a diagram, that a sequence of marks in NRZ-S will produce a constant signal level. What is generated if a sequence of spaces is encoded using NRZ-S?

3. (A) Show, by means of a diagram, that a sequence of spaces in NRZ-I will produce a constant signal level. What is generated if a sequence of marks is encoded using NRZ-I?

4. (A) The illustration of the RS-232 D-connector was said to have female numbering. Find a male RS-232 connector and show how the numbering differs.

5. (A) Read through a PC technical reference manual to determine the electrical signals used within the PC. What is the maximum distance these signals can propagate? (It may be necessary to visit a library and look in an electronics handbook to answer this question.)

6. (A) With an ohm-meter, take an RS-232-C cable and determine if it is a null-modem cable or a "standard" RS-232-C cable. This can be done by holding one of the meter's probes on pin 2 at one end of the cable and the other probe on pin 2 on the other end. If the meter doesn't register any current flow, it can be assumed that pin 2 doesn't connect to pin 2, which means that it is a null-modem cable (verify this by checking pin 2 with pin 3: the meter should register something this time; otherwise, it is a very strange cable). Be careful when testing male connectors; since pins 2 and 3 are adjacent, if the probe is applied to both pins 2 and 3, a false reading on the ohm-meter may be obtained. If the probe cannot make contact with a pin in the female connector, use a paper clip to complete the circuit (unbend the paper clip and place it about 5 millimeters into the pin hole).

7. (B) Determine experimentally what distance an RS-232-C signal can successfully be propagated by connecting a number of RS-232-C cables together. Use the point-to-point telephone utility developed in Chapter 3 to transmit the data from one DTE to the other. Remember to check for the type of cable (i.e., null-modem or standard); if the cables are null-modem, it will be necessary to use odd numbers of cables. Why?

5
Testing Communication Systems

5.1
Introduction

Chapter 3 required working with software that supports point-to-point asynchronous communications between pairs of DTEs; problems may have been encountered in determining

Which DTE was actually transmitting (or receiving).

Which DTE was correctly following the protocol.

Either or both of these problems can result in untold hours of wasted time, trying to determine "which-end-is-doing-what." Not surprisingly, people implementing or maintaining communication systems in the real world also run into these problems, especially when interconnecting DTEs from two different manufacturers, both of which claim to have implemented the same protocol.

Tools are needed that allow the person attempting to solve the problem to monitor the activities on the communication channel, thereby determining exactly which end is transmitting or where the protocol is not being observed. Fortunately, a number of tools are available that satisfy these requirements.

5.2
Built-In Tools

Many UARTs support built-in hardware facilities that permit the programmer to perform numerous software tests. Broadly speaking, these tests fall into two categories: loopback testing and interrupt generation.

5.2.1
Loopback Testing

A loopback test causes all data normally supplied to the UART for transmission to be returned to the processor as if it has just been received. This permits the local software to be tested without errors being introduced by the remote DTE or the communications channel, since the UART is separated from the communication channel. For example, a loopback test makes it possible to test both transmission and reception software locally.

Figure 5.1 demonstrates that when a loopback test is being performed, the serial transmission hardware continues to transmit marks and the serial receive hardware is disconnected.

In the 8250 UART, loopback is set by writing `0x10` to the Modem Control Register (address `0x3FC` for serial port 1 and address `0x2FC` for serial port 2), as in Figure 5.2.

Commkit can be used to demonstrate the UART's loopback feature by changing the Modem Control Register (`MCR_DEFN`, defined in file `devices.h`) to include the constant `LOOPBK`:

```
#define MCR_DEFN    (DTR+RTS+OUT2+LOOPBK)
```

If the above change is made to `devices.h` and the telephone utility described in Chapter 3 is recompiled using make (`commkit.c` will also be recompiled), the UART will be placed in loopback mode the next time `pt2pt.exe` is run. Running the newly compiled version of the telephone utility should cause everything that is typed locally (and appearing in the upper half of the screen) to appear on the lower half of the screen as it would have, had a remote user sent the information. If the modified software does

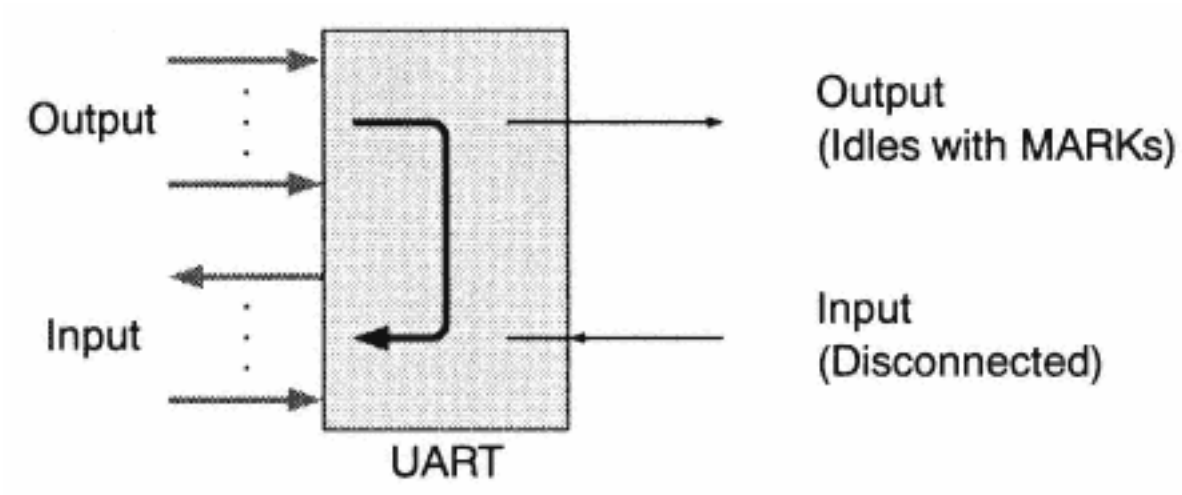

Figure 5.1
The UART loopback test.

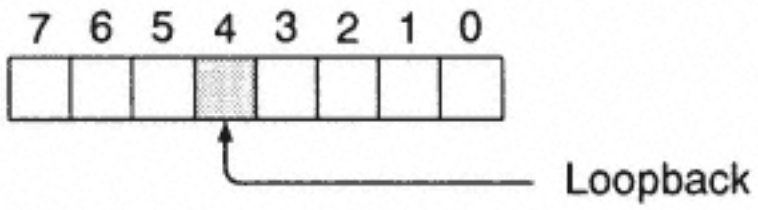

Figure 5.2
Modem Control Register: loopback bit.

nothing (i.e., the lower half of the screen doesn't echo what is typed on the upper half), don't despairnot all 8250 UARTs support loopback. Try running the software on another PC.

The UART is switched out of loopback mode by writing (`DTR+RTS+OUT2`) to the Modem Control Register.

5.2.2
Generating Interrupts

All software, whether a simple first-year assignment or a commercial database package, must be tested before being released. The same holds true for data communications software. However, testing communications software has the added problem that certain errors are caused by random or spurious events that are hard to duplicate. For example, events such as overrun errors and framing errors are unpredictable and may take many hundreds of hours of testing for them to occur (and when they do, they are difficult to duplicate).

Fortunately, the 8250 UART has been designed so that hardware interrupts (such as overrun and framing errors) can be generated through software, meaning that error-handling software can be thoroughly tested before being released.

There are six hardware interrupts that can be software generated by the 8250 from the Line Status Register (Data Available, Overrun Error, Framing Error, Parity Error, Break Interrupt, and Transmit Holding Register Empty), as well as four from the Modem Status Register (the use and testing of the Modem Status Register will be discussed in more detail in Chapter 8). These interrupts are caused by writing to any of the corresponding status bits in the Line Status Register and having the UART in loopback mode; Figure 5.3 shows the specific bits in the Line Status Register.

To illustrate, the steps required to test the overrun-errorhandling software on the second serial port (`SP2`, base address `0x2F8`) are as follows:

1. Turn loopback on (port `0x2FC`).

2. Write `0x02` to the Line Status Register (port `0x2FD`).

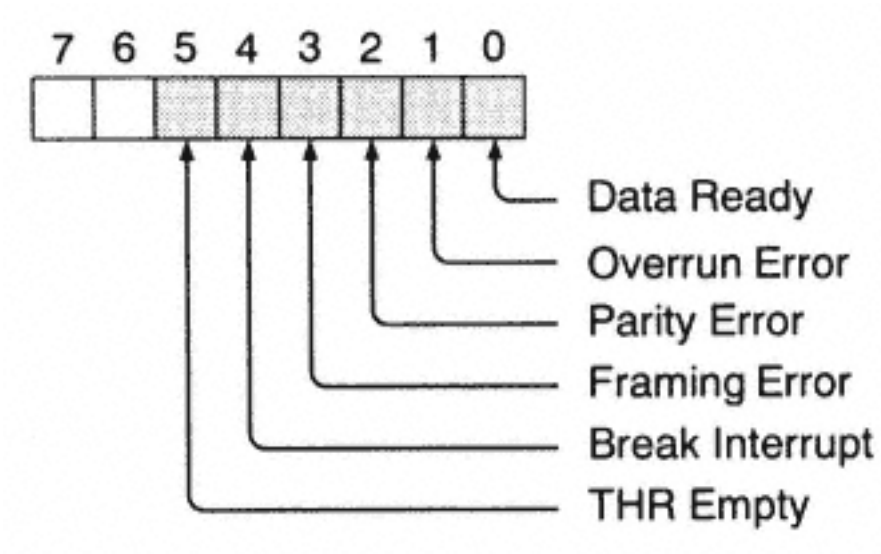

Figure 5.3
Tthe Line Status Register bits
affected in loopback mode.

3. A line status interrupt is generated by the UART, causing control to be passed to the interrupt handler responsible for serial port 2 (`sp2_ih()`).

4. The Interrupt Identification Register will have a value of `0x06` (Receiver Line Status Change) and the Line Status Register will have a value of `0x02`.

5. The interrupt is serviced and cleared as if it were any other interrupt.

Software-generated interrupts can continue as long as the UART remains in loopback mode.

5.3 Channel-Monitoring Tools

Built-in testing can force the hardware to emulate certain conditions for software verification; however, these tests offer no suggestion as to the state of the channel itself. For example, if data transmitted from one DTE fails to arrive at the other, but both DTEs are "working perfectly," something is happening to prevent the data from arriving at the receiving DTE. There are three possibilities:

The data *isn't* being sent by the transmitting DTE.

The data is being sent, but the communication channel is failing to pass the data to the receiving DTE.

The data is being sent from the transmitting DTE across the channel correctly, but isn't being *received* by the receiving DTE.

The first two of these possibilities (and perhaps the third) can be overcome with the capability of observing the activities on the communication channel.

Broadly speaking, there are two types of tool available that permit the observation of the activity between two DTEs on a communication channel: break-out boxes and line analyzers.

5.3.1 The Break-Out Box

In the previous chapter, it was demonstrated that a single communication channel interconnecting two DTEs consists of a number of different connections, each responsible for a specific signal. Since these signals are typically electrical impulses, they can be detected and displayed to the person monitoring the channel.

The break-out box is a tool designed to indicate which signals on each of the different connections between the two DTEs are active. In a break-out box, the state of a signal (i.e., active or inactive) is indicated by a lightemitting diode associated with the connection in question. A break-out box monitoring the signals on an RS-232 communication channel would typically signal the conditions on any or all 25 pins. As data is transmitted between the two devices, for example, the light-emitting diodes displaying pins 2 and

3 become active and inactive, depending upon the value of the transmitted byte (e.g., marks turn the LED on; spaces turn it off).

In addition to monitoring the various signals on the channel, break-out boxes permit testing the channel to activate a nonexisting signal. For example, if pin 2 (transmit) is idling (with marks), the break-out box allows the person testing the channel to force the pin into the space state. Similarly, a signal can be deactivated by breaking the connection.

Break-out boxes generally cost between $50 and $200.

5.3.2 Line Analyzers

Although break-out boxes are useful in determining the state of various signals between the DTEs, they give no indication as to the data that is actually being sent. For example, when a byte is transmitted, the break-out box's transmission LED simply flashes on and off.

Line analyzers (or data analyzers), like break-out boxes, monitor the various signals that occur between the two DTEs. However, line analyzers are designed to *show* the data that is being transmitted on the communication channel rather than only indicating that a signal is present. For example, if the byte *A* is transmitted between a pair of DTEs, a break-out box simply flashes on and off, whereas the line analyzer displays the byte.

At a minimum, a line analyzer permits the user to monitor the data on a communication channel by displaying the transmissions from both DTEs, with the data from each appearing on alternate lines of the display. To allow as much flexibility as possible, line analyzers often have keyboards that permit the user to specify the configuration of the channel (including line speed, parity, and word size); more sophisticated and expensive line analyzers will do this automatically. In addition, some can be programmed to recognize specific protocols and will search for control sequences associated with the protocol before displaying the captured data.

Line analyzers are considerably more expensive than break-out boxes, usually costing anywhere from $3,000 to over $20,000.

5.4 The Commkit Line Analyzer

Commkit is supplied with a software line analyzer for monitoring the communication channel between pairs of DTEs. The Commkit line analyzer can be found in the file `analyzer.c` on the Commkit diskette.

5.4.1 Hardware Requirements

The Commkit line analyzer software requires a PC with two serial ports (`SP1`serial port 1, and `SP2`serial port 2). For example, to monitor the

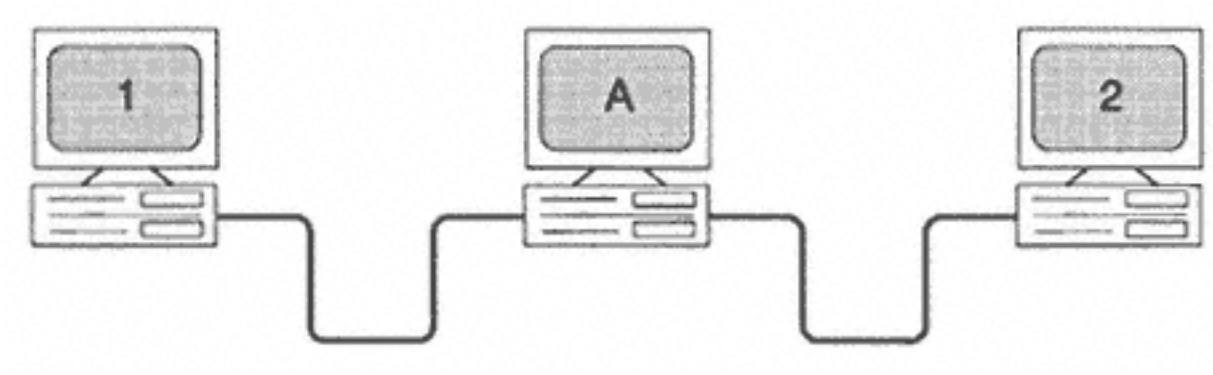

Figure 5.4
The Commkit line analyzer interconnecting two PCs.

traffic between two PCs, a third PC is required to act as the line analyzer. A possible configuration of the three PCs is shown in Figure 5.4 (the analyzer is the PC with the *A* on the screen).

As in all other cases, the PCs must be connected by null-modem cables.

5.4.2
Design

The basic algorithm associated with the Commkit line analyzer is as follows:

1. Read a byte (from either port).

2. Display the byte.

3. Send (forward) the byte onto the "other" port.

The algorithm, as written above, simply displays on the screen each byte received. In order to make the line analyzer output less confusing, each serial port can be associated with a specific set of lines on the analyzer's screen. For example, data received on serial port 1 could be displayed in inverse video on odd-numbered lines (1, 3, 5, etc.), while the data received on serial port 2 could be displayed in normal video on even-numbered lines (2, 4, 6, etc.).

Another useful technique often employed by line analyzers is to interleave the display of the incoming bytes. That is, when a sequence of one or more bytes from one DTE stops, the next sequence of bytes (from either DTE) starts up where the last DTE stopped. This permits the line analyzer to display the events as they occur. For example, in a half-duplex connection, all line analyzer output would be associated with one set of lines (odd or even) until the communication channel was turned around, at which point the second DTE would begin transmitting and the line analyzer's output would appear on the other set of lines (even or odd). However, in a full-duplex connection, output on the line analyzer could appear from either DTE, thereby resulting in the interleaving of output (i.e., a combination of normal and inverse video on alternate lines).

5.4.3
Implementation

The line analyzer is implemented using the message-passing facilities of Commkit. As bytes are received from either of the serial ports on the line

analyzer PC, they are sent in messages to the (foreground) line analyzer process, which writes the byte to the screen and then forwards it to the other port.

The line analyzer software is divided into three distinct parts: the interrupt handlers, the display/control process, and the character-forwarding software.

The Interrupt Handlers

The line analyzer software uses four Commkit interrupt handlers:

`spl_ih()` and `sp2_ih()` . The two serial ports communicate with the analyzer process by making calls to `low_level()` after receiving a character (`RECVDONE`) or after a character has been transmitted (`XMITDONE`). Each character received is forwarded to the analyzer process via queue `APPLICATION`. The transmit done interrupt causes the queue associated with the serial port to be checked for further characters.
When the analyzer process sends a message to either serial port for transmission (code `MSG_AVAIL`), the state of the port is checked. If the port is transmitting, the message is left on the port's queue; otherwise, the queue is read and the byte is transmitted.

`clk_ih()` . The clock interrupt handler calls `low_level()` about 18 times a second; `low_level()` keeps track of the number of ticks. As soon as one second has elapsed, a 1-byte message is sent to the analyzer process via queue `APPLICATION`.

`kb_ih()`. The keyboard interrupt handler calls `low_level()` whenever a character is read from the keyboard. Each character is forwarded to the analyzer process via queue `APPLICATION`.

The Display/Control Process

All bytes received from the serial ports, all keyboard input, and all timing signals are sent to the foreground display/ control process, `do_analyzer()`, via queue `APPLICATION`. `Do_analyzer()` waits for messages and receives them using the `recv()` primitive.

Messages, containing a byte, from either of the serial ports are displayed on the PC's screen (an exception is discussed below) and then forwarded to the "other" serial port queue. The line analyzer process can determine to which port the byte should be forwarded, because the source identifier associated with the byte's message indicates the port from which the byte was received (either `SP1IH` or `SP2IH`).

The variable `base_line` indicates the pair of lines (i.e. `base_line` and `base_line+l`) upon which the output should be displayed. Characters (and their associated attributes, normal and inverse) are written to a specific column within a line (indicated by the variable `column`). Once a character is displayed, `column` is incremented. Should `column exceed` the

screen width, `base_line` is incremented by 2 and `column` is cleared. When `base_line` exceeds the length of the screen (indicated by the constant `END_OF_SCREEN`), `base_line` is set to `START_OF_SCREEN`.

In addition to displaying and forwarding bytes, the line analyzer supports a number of extensions to the original algorithm:

Once a second, a timing signal is displayed on the screen, allowing the person using the line analyzer to obtain an indication of the time taken for each transmission. The timing signal is generated by the clock interrupt handler (`clock_handler()`) and forwarded to the line analyzer process. The timing signal is displayed on the normal video line (i.e. the even-numbered lines) as an inverse-video blank character, and the inverse-video line (i.e. the odd-numbered lines) as a normal video blank character. The timing signal can be toggled (i.e., displayed or not displayed) by pressing CTRL-T (*DC4*), causing the variable `time_display` to be assigned `TRUE` or `FALSE`. By default, the timing signal is displayed when the analyzer is first turned on (i.e., `time_display` is `TRUE`).

Since output from the analyzer can disappear rather rapidly if there is a large volume of traffic flowing between the DTEs, output to the screen can be toggled using CTRL-S (*DC3*) to turn the line analyzer's output off (`FALSE` is assigned to the variable `display_all`), and CTRL-Q (*DC1*) to turn the line analyzer's output on (`True` is assigned to `display_all`). Although transmitted data is not displayed when the line analyzer display is toggled off, the line analyzer still receives and forwards all data. If the line analyzer's output has been toggled off, the timing signal is not displayed.

The line analyzer can be switched off at any time (thereby returning control to MS-DOS) by using CTRL-C (this changes the variable `running` to `FALSE`). At this point, all communications cease since the analyzer is no longer forwarding the characters.

When the line analyzer is started, all data received is displayed (i.e. `display_all` and `time_display` are assigned `TRUE`).

Character-Forwarding Software

Once the character display sequence has finished, `do_analyzer()` sends the character to the other serial port for transmission. However, before the character can be transmitted, the state of the serial port must be determined:

Port is busy. If the port is busy, the character cannot be sent and is left on the queue associated with the serial port. The character (or any intervening characters) are removed one at a time from the queue when each `XMITDONE` interrupt is signaled.

Port is idle. If the port is idle, the character can be removed from the serial port's queue and transmitted immediately using `outportb()`.

The state of the serial port is maintained in the semaphore `transmitting` (an array of two elements, one for each serial port). Each element in `transmitting` is initially 0, indicating that no transmission is taking place; it is incremented whenever a message from the line analyzer process arrives for a specific serial port.

Messages from the line analyzer process cause `low_level()` to be called with the code `MSG_AVAIL`. If the port's `transmitting` semaphore is 0, the queue is read and the character is sent (`transmitting` is incremented in `low_level()` and decremented in `get_char()`). However, if `transmitting` is not 0, `transmitting` is incremented and the character remains queued.

Whenever `XMITDONE` is signaled, the function `get_char()` is called, the queue associated with the interrupting serial port is read, and the character is transmitted. In addition, the serial port's semaphore `transmitting` is decremented.

It was found during some of the line analyzer software testing that once a week (usually on weekends) a serial port would fail to issue an `XMITDONE` signal. The loss of the `XMITDONE` signal stopped all data flow out the serial port because the queue is only read and the data transmitted when `XMITDONE` is detected.

The loss of the `XMITDONE` signal also meant that the serial port failed to read its serial port queue. This failure meant that the pool of message buffers was eventually exhausted (as more messages were forwarded to the unread queue), resulting in the line analyzer crashing and returning control to MS-DOS.

The solution to this problem is to monitor each serial port using a sanity check. The sanity check works on the assumption that as long as `XMITDONE` signals are being generated, `transmitting` should have a value of 0 (not transmitting) or 1 (transmitting). Anything much higher than, say, 10 or 15, indicates that something is wrong with the serial port and an `XMITDONE` signal has been lost. When this point is reached, the sanity check takes over and forces the first character on the queue to be sent by calling `get_char()`. This character then results in the serial port generating an `XMITDONE` signal, meaning that transmissions can resume and the queue will eventually be emptied.

The sanity check occurs whenever a `MSG_AVAIL` is signaled. If `transmitting` exceeds `LIMIT` (defined in `analyzer.c`), it is assumed that an `XMITDONE` is missing and a transmission is forced.

Message Flow

The overall flow of messages within `analyzer.c` is shown in Figure 5.5.

All devices send their messages to queue `APPLICATION`, which is read by `do_analyzer()`. Messages from the keyboard and the clock are pro-

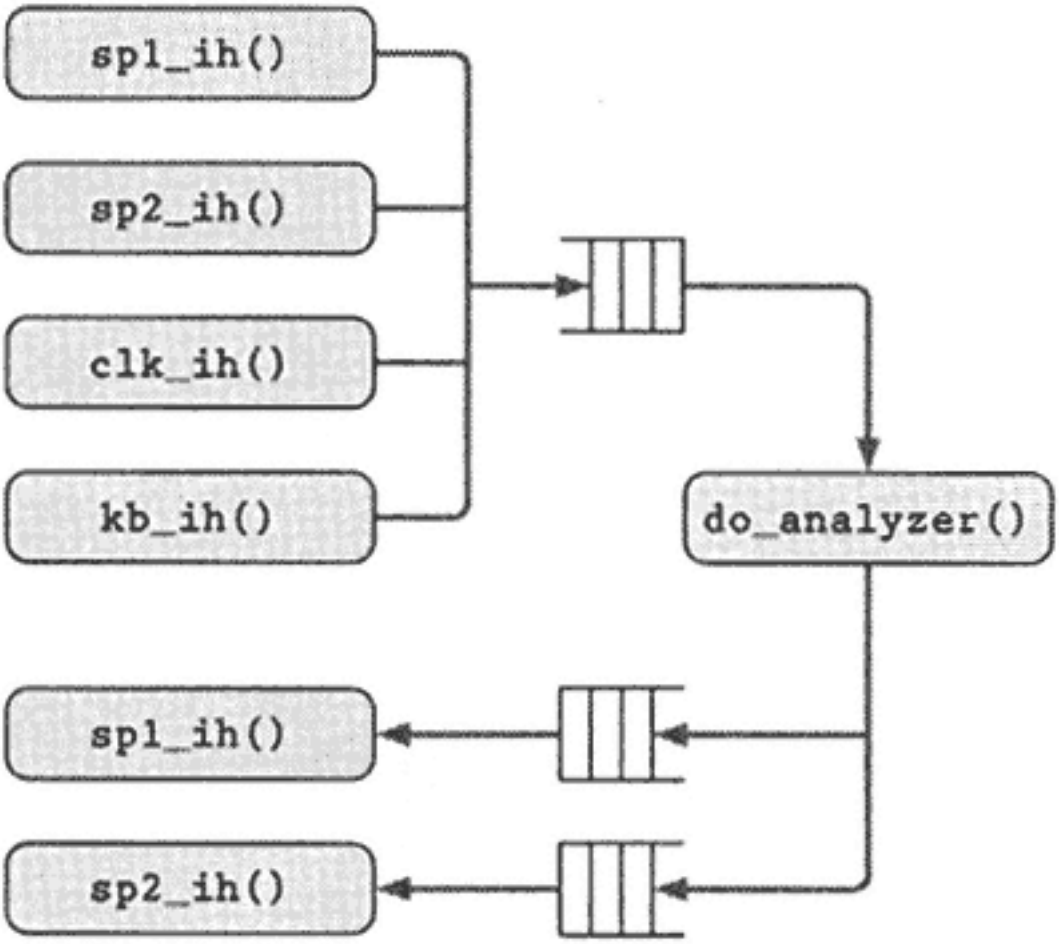

Figure 5.5
Message flow within the Commkit line analyzer.

cessed by the analyzer software; messages from either of the serial ports are displayed (if the display is enabled) and put onto the "outgoing" queue for transmission. All messages are a maximum of one byte in length.

5.5
Using the Commkit Line Analyzer

The Commkit diskette is supplied with both the source and the executable code for the line analyzer. The line analyzer requires a PC with two serial ports.

5.5.1
Creating a Commkit Line Analyzer

An executable line analyzer is made from the line analyzer source code, `analyzer.c`, `commkit.obj`, `srcalls.obj`, and the standard header files (`general.h`, `ascii.h`, and `devices.h`):

```
C:\> make analyzer.exe
```

5.5.2
Running the Commkit Line Analyzer

Once compiled, the line analyzer can be executed by typing `analyzer` and the required line speed (50, 300, 1200, 2400, 4800, or 9600) after the MS-DOS prompt. For example, to run the line analyzer at 2400 bits per second, one types:

```
C:\> analyzer 2400
```

When using the line analyzer, remember to start the line analyzer software (on the analyzer PC) *before* starting any of the other PCs. This ensures that all data transmitted between the two PCs will be forwarded rather than blocked by the inactive line analyzer PC. Should the line analyzer be shut down before a transmission has completed, the remaining bytes will not be forwarded.

The control characters recognized by the line analyzer are as follows:

CTRL-Cterminates the line analyzer, returning control to MS-DOS.

CTRL-Tturns the 1-second timing pulse on or off.

CTRL-Sstops display of all output (data and timing signals).

CTRL-Qresumes display of all output (data and, if enabled, timing signals).

5.6 Summary

This chapter examined some of the tools available to assist in the testing and development of communications software. Some rudimentary tools are available within the UART itself, allowing conditions such as overrun and the break indication to be tested directly by the processor. As useful as these tools are, though, they do not offer the user any indication as to what is happening on the communication channel itself.

Two different tools are available for examining the traffic on a channel: the break-out box and the line analyzer. Break-out boxes permit the monitoring of various signals on the channel, such as whether a byte has been sent or received. However, no indication is given as to the value of the data. Line analyzers, on the other hand, allow the user to actually see what is happening on the channel, aiding in the development and implementation of communications software.

Commkit is supplied with a line analyzer that can monitor the traffic flow between two DTEs. This analyzer offers some of the features found on commercial analyzers, at a somewhat reduced cost.

Finally, it is important to remember that there are three possible places to consider when testing communication software: the transmitter, the channel, and the receiver. Tools such as the line analyzer allow a software developer to monitor the traffic on a channel and help in tracking down software errors.

5.7 Exercises

1. (C) Run a loopback test on the telephone utility described in Chapter 3.

2. (B) Run the PC phone utility (developed in Chapter 3) with the line analyzer connecting the two PCs. Try typing a message from one PCwhat appears on the analyzer? Use the timing signal to determine how long it takes to type an entire line of characters. Type messages from both

PCswhat appears on the line analyzer? Toggle the line analyzer so that output does not appear on the analyzer's screen; now type messages on both PCsdoes the information still get through?

3. (C) Modify the line analyzer to allow the user to see the data displayed in hexadecimal, octal, or ASCII. Let the user enter the selection from the keyboard while the program is running (permitting the user to change the output dynamically).

Test the program by modifying the PC phone utility so that it sends 7 bits of data and parity (mark, space, odd, or even). The line analyzer should read eight data bits and *no* parity. The value displayed by the line analyzer will be the data and the parity bits.

(Note that for hexadecimal and octal output, a single character will map into 2 or 3 display characters.)

4. (C) As mentioned in Section 5.4, some of the more sophisticated line analyzers can dynamically configure themselves by monitoring the traffic on the line (i.e., without operator intervention). Modify the Commkit line analyzer so that it will configure itself to the speed of the line and the word size.

In order to do this exercise it will be necessary to consider the different line status signals generated by the serial ports. Instead of recording the error statistics in `port_handler(), low_level()` will have to be called with the value of the Line Status Register.

5. (C) Add a clock facility to the line analyzer that will display a clock in the upper right corner of the screen. The clock can be useful in a number of situations; for example, timing the length of a transmission. Use CTRL-R (*DC2*) to reset the clock to 0.

6. (C) Add a scanning feature to the line analyzer. That is, allow it to search for specific sequences of bytes. Prior to finding the search string, no data should be displayed; however, once the search sequence is found, all subsequent data received should be displayed. Allow the user to change the sequence of characters being scanned while the analyzer is running.

By default, the analyzer should not search for any strings. Use CTRL-F (*ACK*) to signal that the user wants to enter a search string (the number of characters in the string should not exceed 32). Once a string has been entered, it should not be searched for until the user types CTRL-U (*NAK*), at which point a message should appear on the top line of the analyzer's screen indicating that a search is under way. Typing CTRL-U again stops the search.

In many situations it is useful to search for both the starting and ending strings in a sequence of characters. Extend the analyzer so that the user can specify an "end-string": the end-string should be accepted after the user types CTRL-E (*ENQ*). The end-string should not be searched for until the starting string has been found. As soon as the end-string has been found,

the starting string should be scanned for once again and no data displayed until it has been found.

All data should be forwarded while the search is taking place.

7. (C) Use the line analyzer to determine the terminal control characters sent between a local host and a terminal connected to the host. Once these characters have been obtained, modify the point-to-point software to emulate the control characters sent from the host. Use the line analyzer to debug the terminal emulation program.

Although the easier way of solving this problem is with the terminal's manual, if the manual is not available, this technique can be used to determine the control characters.

PART 3
LONGER-DISTANCE COMMUNICATIONS

Consider the following situations:

A file containing thousands of binary records (perhaps millions of bytes) must be transferred from one computer to another without the loss or corruption of any information. How long will the transfer take, and can it take place using only asynchronous point-to-point communications?

Imagine working in the branch of an office that is several hundred miles from the head office, where the central computer facility happens to be located. How will the information stored at the head office be accessed?

Transferring large volumes of binary information using the point-to-point techniques discussed in Chapter 3 are both potentially error prone and time consuming. Consider, for example, how the transmitting computer can determine whether the destination computer had received a byte correctly. In a terminalhost asynchronous configuration, the error can be checked by remote echoing, in which the character typed by the user is echoed by the remote host. If the character echoed is not the same as the character entered, the user can delete the character in error and type the right one. The same approach does not lend itself to the transfer of large volumes of data for a number of reasons:

Checking by echoing each character is a costly, time-consuming operation, since the throughput is essentially halved with all the characters being echoed.

Simply that a character is echoed does not mean that it was received correctly. In the worst case, the character may be corrupted before being received (for example, from *A* to *C*) and then corrupted again as it is echoed

(for example, from *C* to *A*). In this situation, neither the transmitter nor the receiver has any way of recognizing that the character was received incorrectly.

If the character echoed does not match the one originally sent, the transmitter can detect the error, but how can it be corrected? Retransmitting the character is of no use, since the receiver has had no indication that there was an error with the previous character.

If binary data is being sent, a delete character must not be interpreted as a delete operation but rather as part of the data being sent.

There is a second problem to consider when dealing with the transfer of large volumes of data using asynchronous communications: the overheads associated with the transmission. When an asynchronous frame format is used to transmit a byte of information, the number of control bits range from a minimum of two (if only one start bit and one stop bit are used) to a maximum of four (if the start, parity, and two stop bits are used); in the worst case, when sending a 7-bit byte (with the four control bits) a 36-percent overhead can be expected. For example, if the file in question contains 1 million bits, a total of over 1.36 million bits must be sent in the transmission. This is a potentially costly overhead if the user is expected to pay connection charges.

Problems also arise when attempting to use RS-232-C (or any of the other standards that were discussed in Chapter 4) over large distances. For example, RS-232-C signals are attenuated once the distance between the DTEs reaches over about 100 meters. This means that using RS-232-C to transmit over any reasonable distance (between buildings or between continents) is impossible without some form of assistance.

Part 3 presents a step-by-step examination of the techniques used to overcome the limitations of distance, volume, and error handling in data communications. Methods of increasing the amount of *useful* information transmitted (by decreasing the amount of framing information associated with each byte) using synchronous communications are discussed in Chapter 6. Chapter 6 also presents a number of error detection techniques that far surpass the simple parity check.

Chapter 7 introduces some of the concepts associated with transferring files between DTEs. The chapter examines problems such as reliable file transfer and how the file system can interact with the communication system.

The remaining three chapters look at methods of overcoming the distance limitations imposed by standards such as RS-232-C. Chapter 8 shows how the most ubiquitous communication system of all, the telephone, is used to support data communications. Using the telephone system for data communications can prove to be an expensive solution, especially when the communicating DTEs are separated by distances that incur charges for the call. In Chapter 9, a number of multiplexing techniques are presented that can be used to share a single communication channel among a number of DTEsreducing the separate physical communication channels required.

Part 3 concludes with a discussion of why using the existing telephone system for data communications is not necessarily a good thing and considers alternatives. Chapter 10 shows how ideas associated with synchronous communications, the telephone system, and multiplexing all came together in the early 1970s to allow the development of long-distance communication facilities entirely devoted to data communications, known as wide area networks.

6
Synchronous Communications

6.1
Introduction

Asynchronous communications are intended for low-volume data transfers, consisting of several bytes a second, typically between a user typing at a local DTE connected to a remote DTE. However, if the transfer between the two DTEs is a large-volume transfer, consisting of thousands or even millions of bytes (as could easily happen if a file is being transferred), asynchronous communications may prove to be a less desirable data transfer method. Consider the following:

Transmission overheads. Each byte that is sent requires at least two, and possibly as many as four, extra control bits. These control bits are necessary, since they ensure that the byte is properly framed (the start and stop bits) and offer a degree of error detection (the parity bit). With the inclusion of these control bits, each asynchronous frame transmitted has an overhead of 20 percent (a 10-bit frame consisting of an 8-bit byte and 2 control bits: start and stop) to over 36 percent (an 11-bit frame consisting of a 7-bit byte and 4 control bits: start, parity, and two stop). Therefore, for every million bits sent, between 200,000 and 360,000 bits are transmitted as control information.

Error detection and correction. The only means available to the receiving DTE to determine whether the byte has been received correctly is the parity bit. The parity bit can help in detecting errors, but can only detect an odd number of bit changes. It is adequate for low-volume transfers in which each byte is echoed by the remote DTE, since the user can tell immediately whether what has just been typed is echoed correctly. However, echoing each byte in a high-volume transfer halves the throughput.

One way in which DTEs can transfer less control information and more data is to use synchronous communications. Synchronous communications are intended to

Minimize transmission overheads by reducing the amount of control information sent with each message (i.e., sequence of data bytes).

Support better error handling, thereby making error detection and correction more efficient.

The remainder of this chapter examines how the above two goals can be achieved using synchronous communications.

6.2
Reducing Transmission Overheads

Ideally, when a message is sent between a pair of DTEs, nothing other than data is transmitted. For example, in Figure 6.1, three data bytes are transmitted.

However, it was shown in Chapter 3 that two problems must be overcome if data is to be sent between DTEs:

Frame delimiting. The channel idles in the mark state, and any byte sent with one or more leading marks (1s) will not be detected correctly by the receiving DTE.

Timing. The transmitting and receiving DTEs must remain in step to ensure that each bit is received as transmitted.

In an asynchronous communication, these problems are overcome through the use of additional control bits (the start and stop bits) and by restricting the number of bits that can be transmitted (the byte size plus the control bits).

Since synchronous communications permit the transmission of data without embedded control characters, techniques must exist to overcome the problems of frame delimiting and timing.

6.2.1
Frame Delimiting

In an asynchronous communication, all frames are prefixed with a start bit to distinguish the data from the channel's idle state. A similar solution can be applied to sending a message using synchronous communications: before

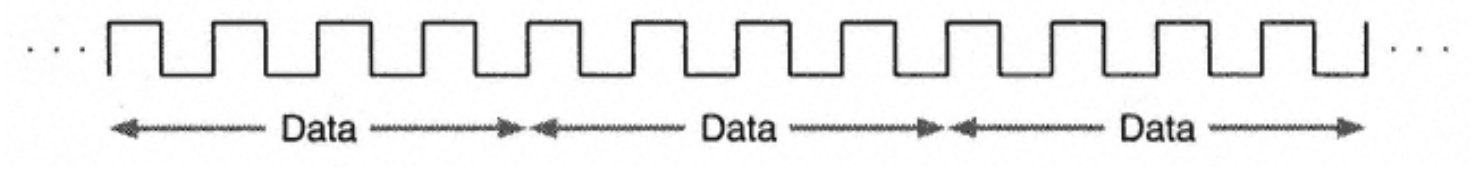

Figure 6.1
Transmitting "pure" data.

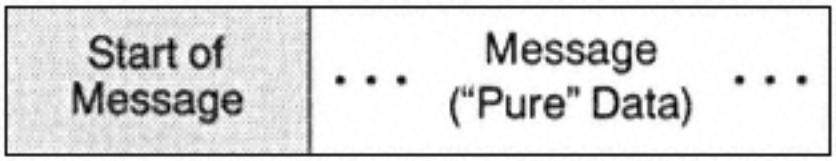

Figure 6.2
Prefacing a message with a start-of-message indicator.

the first data byte is sent, the transmitting DTE forwards an indication (such as a special byte) that a message is about to follow. Figure 6.2 is an example (note that in this and subsequent examples, all transmissions occur from *left* to *right*).

Upon receipt of the start-of-message indicator, the receiving DTE can begin removing the message from the channel. If the start-of-message indicator is a byte, there is no need to support special hardware for its transmission and reception.

A similar argument can be applied to ending the message: the receiving DTE must be able to detect the end-of-message; otherwise, it will continue receiving indefinitely. Three ways in which a transmitter can indicate to the receiver where the end-of-message occurs include:

The first, and simplest, is to have all messages consist of the same number of bytes. By using a fixed message size, the receiving DTE always expects a message to consist of, say, *N* data bytes. Once the start-of-message indicator is detected, the receiving DTE samples the channel for *N* bytes of data. The transmitting DTE always sends *N* bytes after the start-of-message indicator.

The problem with using a fixed message size is that the message to be sent may consist of *fewer* than *N* bytes, requiring the message to be padded with extra bytes that are not part of the message itself. A mechanism must be available whereby the transmitting DTE can inform the receiving DTE of the number of data bytes actually sent.

A second message-delimiting technique involves the transmitter sending the receiver a byte containing a count of the number of bytes in the message. The byte count is usually sent immediately after the start-of-message indicator, as shown in Figure 6.3.

The byte count can be used by the receiving DTE to count the number of incoming bytes. When the count reaches 0, the receiving DTE can stop receiving:

```
receive(Start_Byte);
receive(Count_Byte);

for (i = 0; i != Count_Byte; i++)
     receive(Message[i]);
```

Although the algorithm for using a byte count is quite simple, the drawback

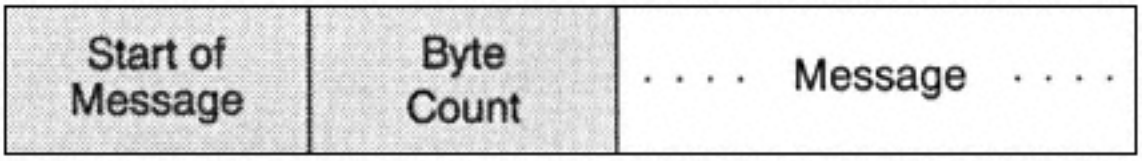

Figure 6.3
Prefacing a message with a byte count.

is that if the byte count is incorrect (for example, its value may have been changed because of a noisy channel), the receiver will receive too many or too few bytes.

A third approach is to transmit a special end-of-message indicator after the last byte of the message. The incoming byte stream is always scanned for the end-of-message indicator; once it has been detected, the receiving DTE has received the entire message. For example:

```
receive(Start_Byte);
receive(Data_Byte);
i = 0;
while (Data_Byte != END_OF_DATA)
{
      Message[i++] = Data_Byte;
      receive(Data_Byte);
}
```

Messages sent with byte counts often have an end-of-message indicator appended to the last byte of the message.

The start-of-messagemessageend-of-message sequence is known as a frame and is shown in Figure 6.4.

Should the end-of-message indicator be lost (for example, because of noise), the receiving DTE will receive more bytes than originally sent. The solution (which can also be used in the byte count method) is to limit the number of bytes that can be sent in a frame (and hence received). As an example, assume that a frame cannot contain more than *N* bytes and that the frame is delimited by the bytes `START` and `END`. The receive data algorithm can be written as follows:

1. Wait for the `START` byte.

2. Set `byte_count` to 0.

3. Check each data byte as it arrives:

`END` byte: Message with valid number of bytes received. Terminate loop. `default`: Check `byte_count`: `0 .. N-1:`
Store byte in buffer, increase `byte_count` by 1,repeat from step 3. `N`: Error, missing `END` byte. Terminate reception process.

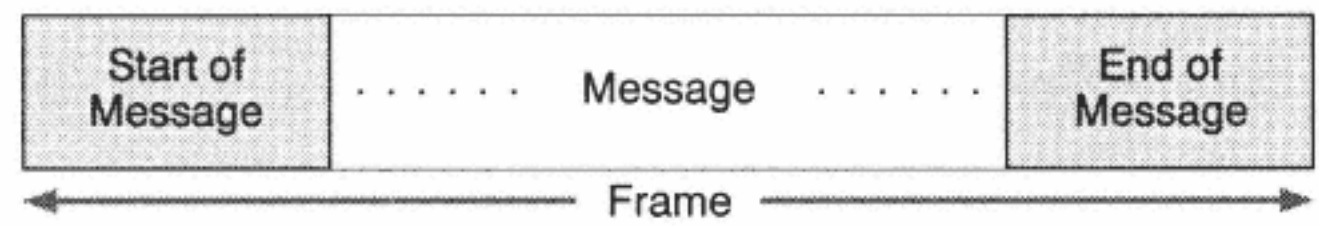

Figure 6.4
A frame encapsulating a message.

The above algorithm ensures that a receiving DTE will not continue receiving indefinitely (and possibly run out of buffer space).

In an asynchronous communication, a frame consists of a single byte. The number of bytes in a synchronous frame depends upon several factors:

The number of bytes needed to represent the control sequences and other framing information. Ideally, this is very small in relation to the total frame size.

The amount of storage available at the receiving DTE. The data bytes within the frame must not exceed the receiving DTE's storage capacity.

The amount of information that must be retransmitted should an error be detected in a frame. For example, if a file of 1 million bytes is transmitted as a single frame, and if one byte is found to be in error, the entire 1 million must be retransmitted. However, transmitting the file as a series of 1,000-byte frames means that the amount of information to be retransmitted is considerably less.

The time required to transmit a frame, which must not exceed the length of time the hardware can remain in synchronization. If synchronization is lost, the contents of the frame will be lost, potentially requiring another transmission (preferably using a smaller frame size).

6.2.2
Timing

Information that is transmitted serially between DTEs is transmitted a bit at a time. To achieve a high probability that each transmitted bit is received correctly, both the transmitting DTE and receiving DTE must be synchronized.

In Chapter 3, it was shown that a UART's clock is designed to remain in step with the remote UART only for the time it takes to send a single byte. The transmitter's and receiver's clocks are synchronized by the start bit: once the center of the start bit is found, the channel can be read once every bit-time. As soon as the receiving DTE has assembled the byte, the two DTEs are no longer synchronized, meaning that the next byte sent requires that a start bit be prefixed.

If "pure" data is sent between the communicating DTEs, the DTEs must still remain in step to ensure that each byte sent is the same as the byte received. However, a synchronous communication, unlike an asynchronous

communication, requires that the DTEs remain in step for tens, hundreds, or possibly even thousands of bytes. Keeping the DTEs synchronized is not a trivial matter, since clocks can drift and result in the loss of synchronization.

As with so many other aspects of data communications, a number of techniques are available that permit the DTEs to transmit a frame of "pure" data and at the same time stay in step:

Special synchronization sequences. Special synchronization byte sequences can be scattered throughout the frame. These characters are detected by the hardware of the receiving DTE, which can realign its clock to correct for any drift that may have occurred. As shown in Figure 6.5, frames are often prefixed by the synchronization sequences to ensure that the clocks are properly aligned before the transmission begins.

Since the synchronization sequence is being embedded within the frame of "pure" data, it is necessary to distinguish between synchronization sequences used to realign the clocks and data bytes that happen to have the same value as that of the synchronization sequence (an event that can occur if a binary file is transmitted). Protocols have been designed to overcome this problem and are discussed in subsequent sections.

Separate timing channel. Another solution to the problems associated with the synchronization of DTEs is to transmit on two channels: one reserved for the data bits, the other for a clocking signal. With this method, as illustrated in Figure 6.6, the receiving DTE knows exactly when to sample the channel since the transmitting DTE indicates the center of the data bit with a timing pulse.

There is another reason for having one of the DTEs supply the clocking signal: the other DTE may not support a clock:

If the transmitting DTE is supplied with the clock, it sends each data bit on one channel and a timing pulse on the other. The receiving DTE

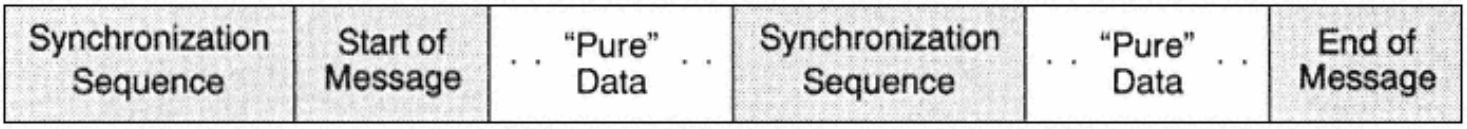

Figure 6.5
Achieving synchronization using synchronization sequences.

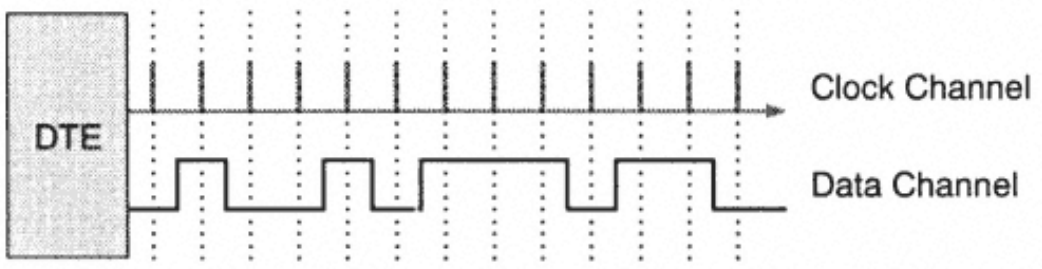

Figure 6.6
Data transmission using a separate timing channel.

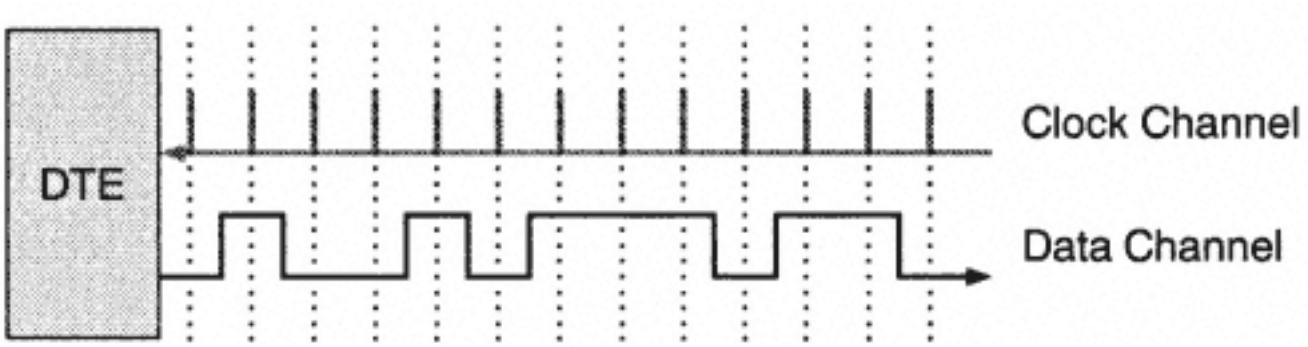

Figure 6.7
Data transmission using an external clock channel.

detects the timing pulse and reads the data channel to determine the value of the bit.

If the transmitting DTE is not supplied with a clock, the receiving DTE must supply the timing signal. In this situation, the transmitting DTE waits for a timing pulse and then sends a bit; the receiving DTE can sample the channel after the timing pulse has been sent, as shown in Figure 6.7.

Encoding the timing into the data. A third approach to keeping the communicating DTEs synchronized is to encode the clocking information directly into each bit as it is transmitted. In this way, as the signal arrives, the receiving DTE extracts both the value of the bit and the clock signal.

One of the best known examples of encoding the clock into the data stream is Manchester encoding, in which the middle of each bit contains the clock signal as well as an indication of the value of the bit (see Figure 6.8).

In Manchester encoding, a low-to-high transition in the middle of the bit interval means a space (or 0), while a high-to-low transition indicates a mark (or 1). Figure 6.9 shows the encoding of the byte *S* using Manchester encoding.

The extracted data is half a signal behind the transmitted data because the value of each bit is determined at the point where the clock signal is extracted (i.e., the center of the bit).

6.3
Improving Error Detection

In a full-duplex asynchronous communication, error detection is typically the responsibility of the user. In remote echoing, if the character typed is not the

Figure 6.8
Manchester encoding.

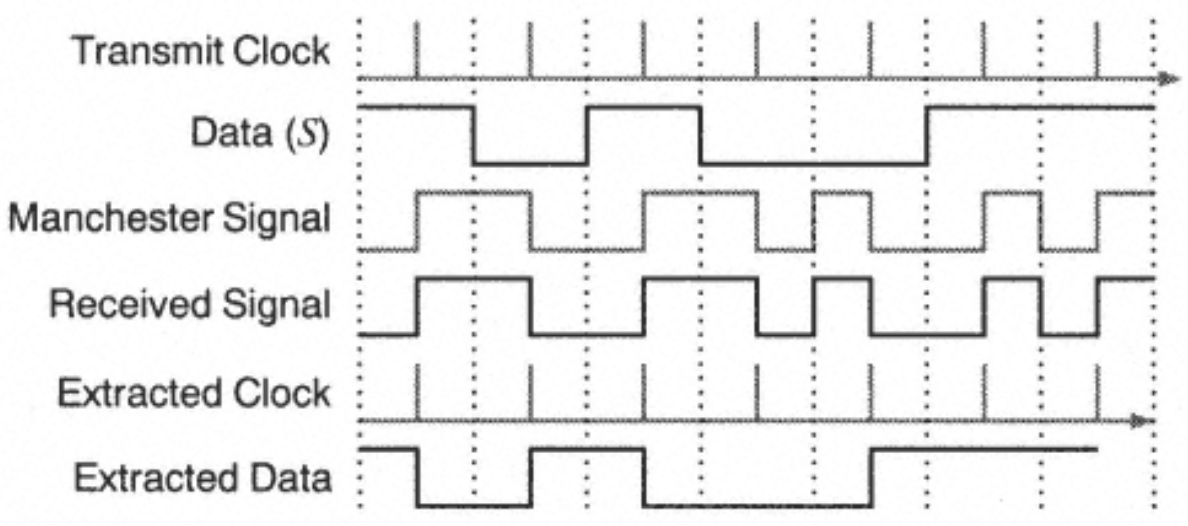

Figure 6.9
The Manchester encoding of the byte *S*.

same as the character echoed, the user can deduce that an error has occurred and that the character must be retyped. Should the receiving DTE detect a parity error, an error message can be returned to the user, requesting that the character be re-entered.

However, as Figure 6.10 illustrates, even if each byte is echoed when received, there is no guarantee that the transmitting DTE can determine that an error has occurred simply by checking the echoed byte and its parity.

Should the transmitting DTE determine that an error has occurred, it must inform the receiving DTE that the last byte sent was in error and that the correct byte follows. For example, the character % can be sent by the transmitter when the echoed character is not the same as that originally sent; a possible scenario is shown in Figure 6.11. However, this solution simply leads to other problems:

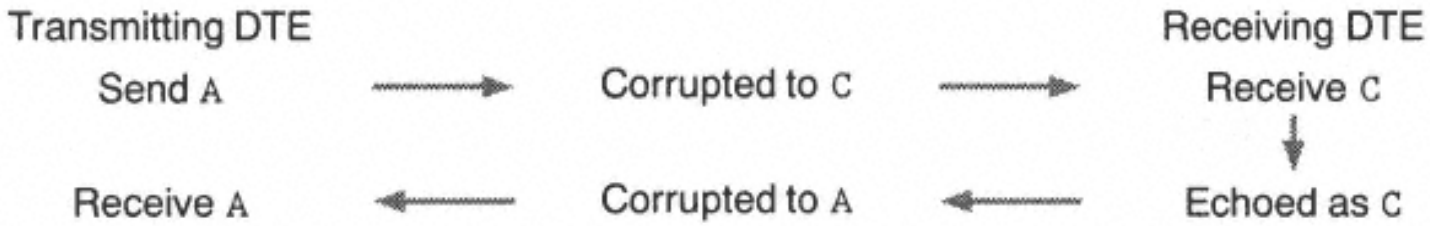

Figure 6.10
A limitation of error detection using echoing.

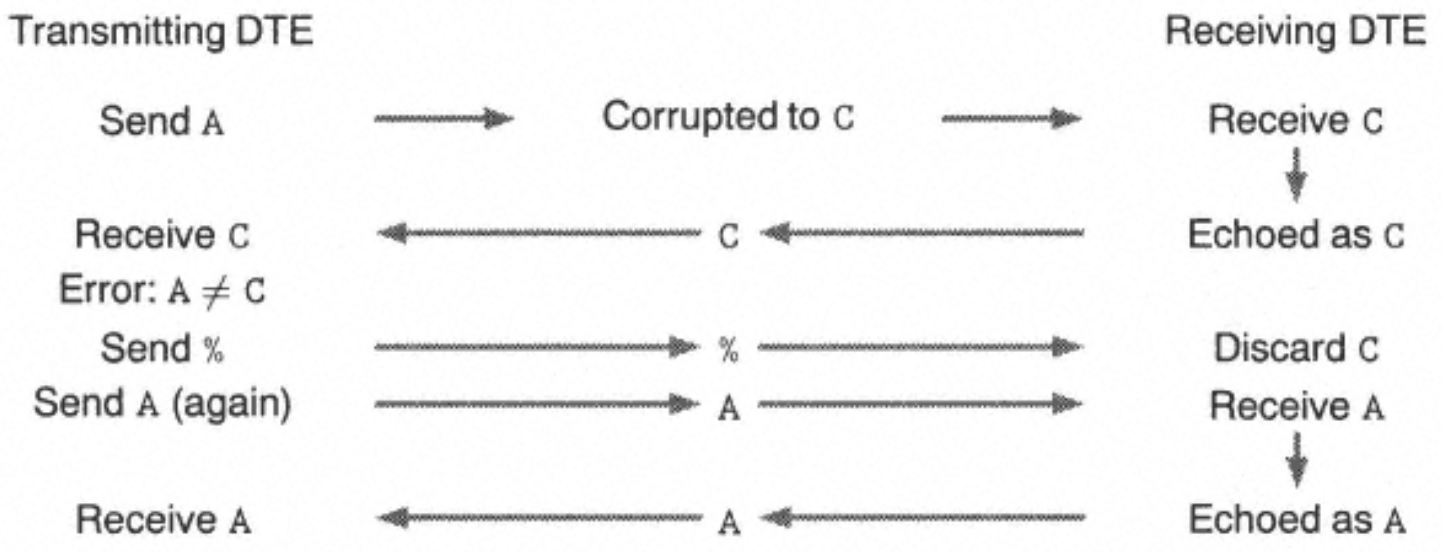

Figure 6.11
An attempt at error correction using echoing.

A mechanism must be devised whereby it is possible to send a % as data (that is, the receiving DTE must be able to distinguish a data % from a retransmission request %).

If an error occurs during a retransmission (especially if the % is corrupted), a protocol must be devised that allows the transmitter to instruct the receiver to discard the corrupted %.

If for some reason the % is lost, the receiver will accept the retransmitted byte but not discard the original corrupted byte.

Admittedly, it is possible to write software to handle the above and other problems, but in the end the overheads (such as echoing each byte) will make the solution unnecessarily clumsy and expensive. In addition, the parity check is costly in terms of the amount of information that must be sent and the level of detection achieved. For example, if a message of 200 bytes is sent, 200 parity bits are also sent.

Fortunately, a number of techniques exist that are more accurate in their ability to detect errors and less costly in terms of the amount of information transmitted when applied to large volumes of data.

6.3.1 Longitudinal Redundancy Check (LRC)

The parity check is used to check an individual byte by summing the bits *across* the byte. A second type of parity check, known as a longitudinal redundancy check, or LRC, can be applied to a "column" of bits within a message. That is, whereas parity is applied to a single byte in the message, the LRC is applied to *all* of them.

The LRC is a byte, initially set to `0xFF`. Prior to transmission, each data byte is exclusive-or'ed (XOR) with the LRC. (The output from an exclusiveor is 0 if the bits are the same and 1 if the bits are different.)

Once all of the data bytes have been transmitted, the LRC is sent. Consider the transmission of a message (`msg`) of length *N*:

```
LRC = 0xFF;
for (i = 0; i < N; i++)
{
    LRC ^= msg[i];  /* Equivalent to: LRC =
                       LRC ^ msg[i];  ^ denotes XOR */
    xmit (msg [i]);
}
xmit(LRC);
```

For example, if the above algorithm were applied to the following five bytes:

```
01101100
10101111
```

```
01110101
11100010
00010111
```

the LRC result would be as shown in Figure 6.12.

The receiver follows a similar algorithm, initially setting the LRC to `0xFF`. As each byte is received, it is exclusive-or'ed into the receiver's LRC:

```
LRC = 0xFF;
for(i = 0; i < N; i++)
{
      recv(&msg[i]);
   LRC ^= msg[i];
}
recv(&last);
LRC ^= last;
```

Assuming that no errors have occurred during the communication, the transmitter's LRC will be the same as the receiver's, meaning that when the receiver exclusive-or's the transmitter's LRC, the result should be 0. Figure 6.13 shows the results of transmitting the five data bytes followed by the LRC byte.

When the last byte is received and the LRC is 0, it means that no errors were detected. Since the LRC algorithm, like that of the parity check, uses exclusive-or'ing to calculate the check information, it should not be surprising that the LRC is no more accurate than the parity check. For example, if an odd number of bits within a column are inverted because of noise, the error can be detected; however, an even number of bit changes cannot be. This problem is illustrated in Figure 6.14.

In some manuals, the term vertical redundancy check, or VRC, can be found in addition to "LRC." "VRC" is simply another term for parity check. Both "VRC" and "LRC" refer to the error checking used with magnetic computer tapes. In tape lexicon, "vertical" refers to information stored across the width of the tape (i.e., data bytes stored in parallel with a parity bit), while "longitudinal" refers to information stored along the length. In Figure 6.15, a vertical redundancy check is applied across the tape and a longitudinal redundancy check is applied along the tape's length.

0	1	1	0	1	1	0	0	Byte 1
1	0	1	0	1	1	1	1	Byte 2
0	1	1	1	0	1	0	1	Byte 3
1	1	1	0	0	0	1	0	Byte 4
0	0	0	1	0	1	1	1	Byte 5
1	0	1	1	1	1	0	0	Check Character (LRC)

Figure 6.12
The LRC after XOR'ing five data bytes.

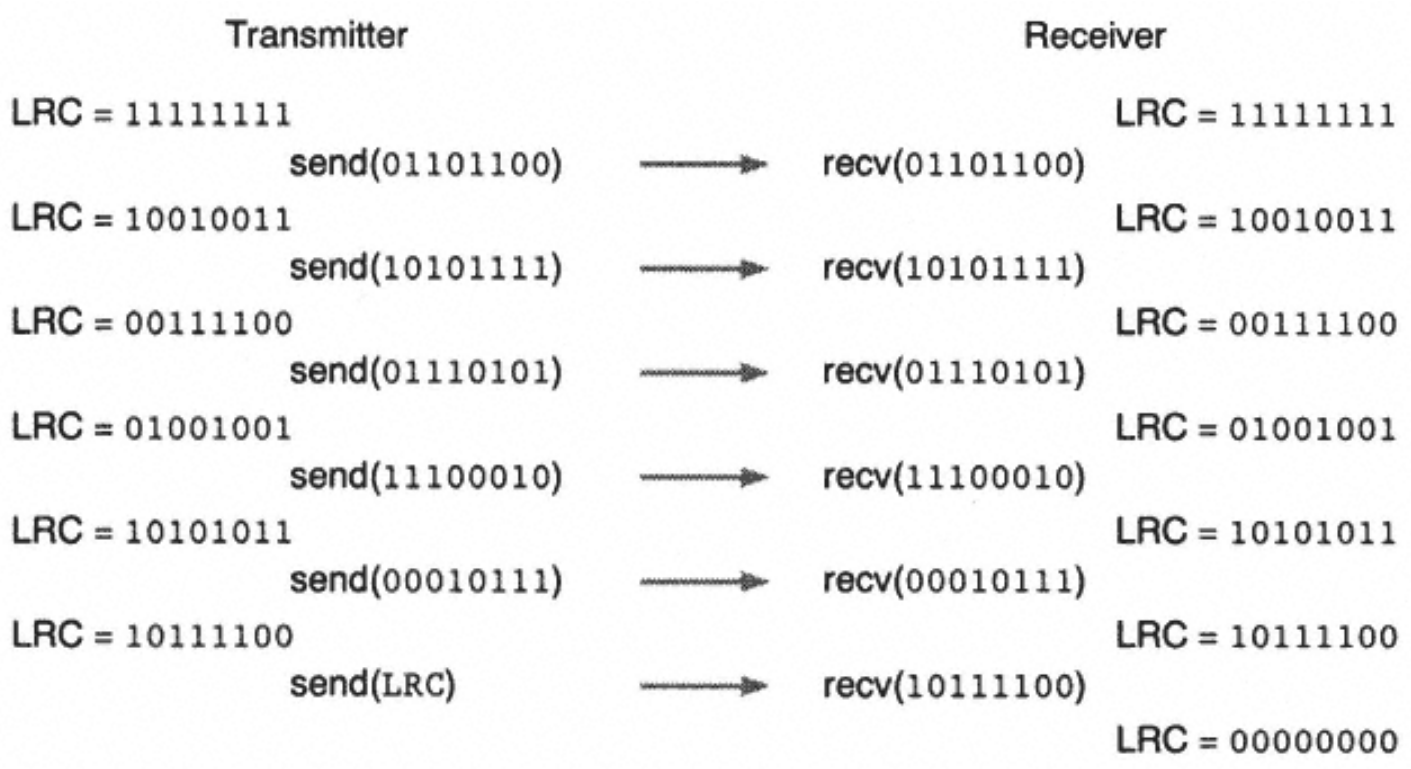

Figure 6.13
The transmission of a message and its LRC.

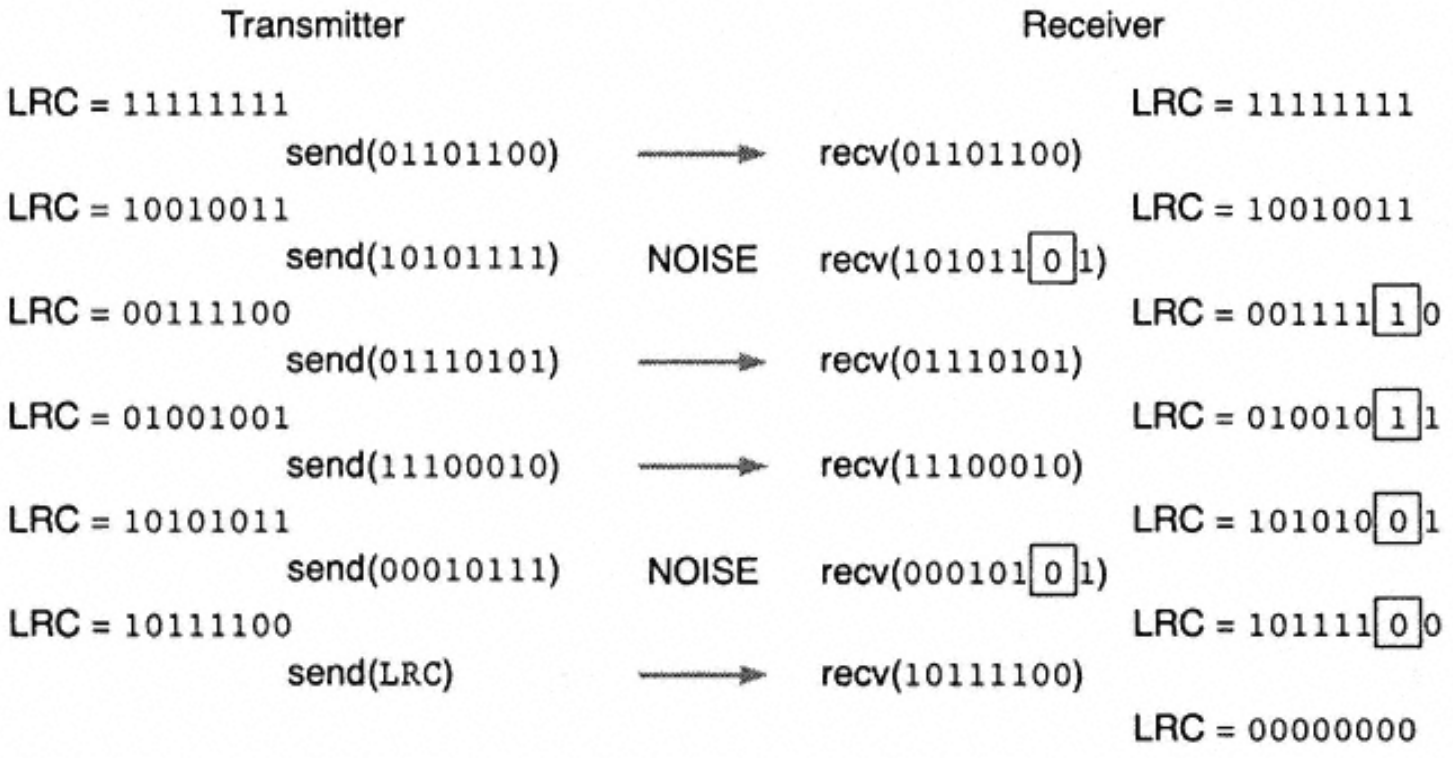

Figure 6.14
An error that is undetectable by the LRC.

Error detection can increase by two to four orders of magnitude over either VRC or LRC alone if the two techniques are used together (that is, each byte is sent with a VRC and is included in an LRC). In Figure 6.16, five 8-bit bytes are sent (with hardware-generated *odd* parity), followed by the LRC.

Although the LRC includes all the data bits, it does *not* include the hardware-generated parity check bit. In the above example, the LRC's parity check is 0 (because the LRC is sent as odd parity and has an odd number of bits set) rather than 1 (as it would be if the parity bits were included in the LRC).

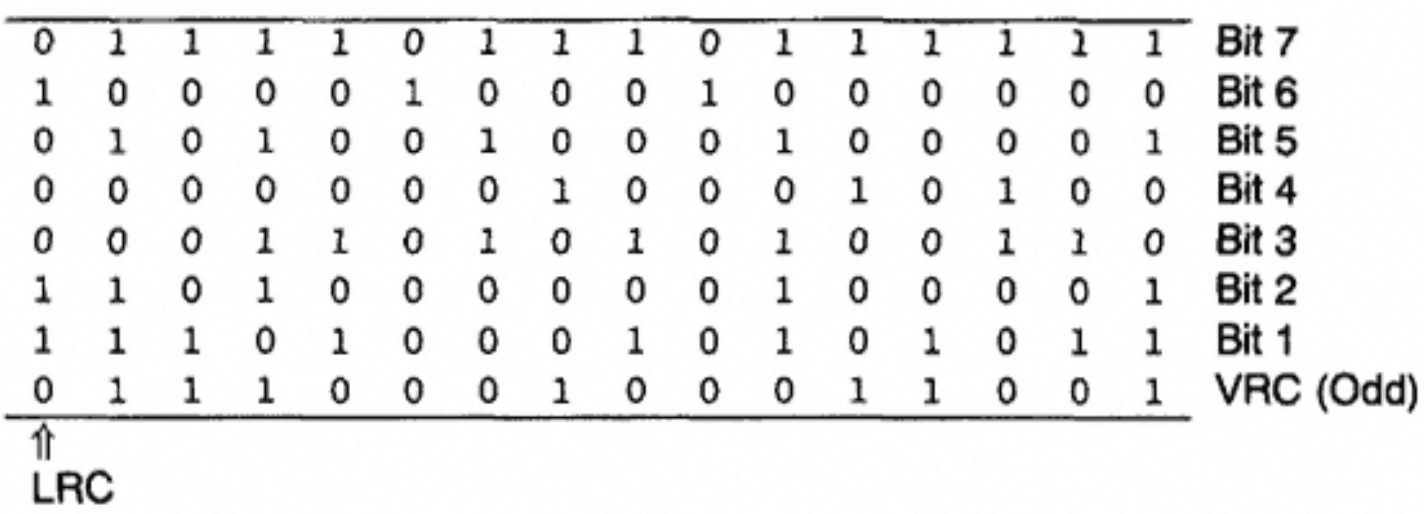

Figure 6.15
VRC and LRC encoding.

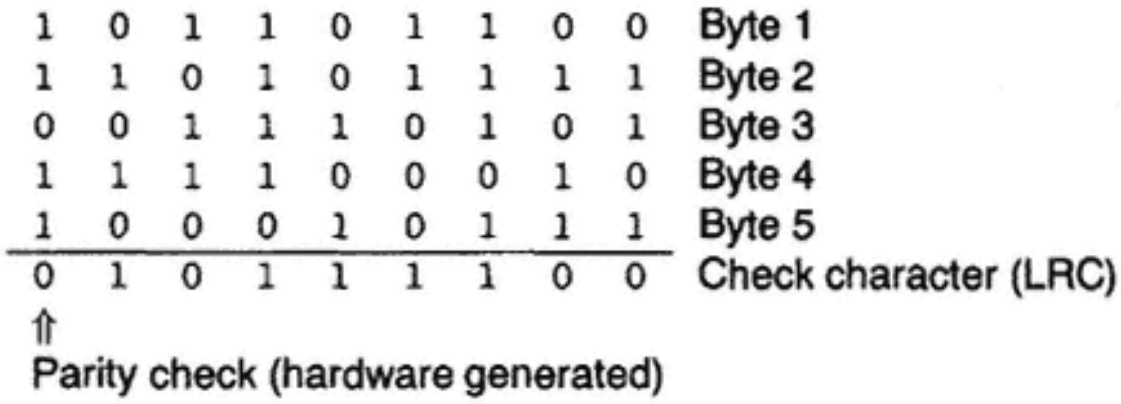
1 0 1 1 0 1 1 0 0 Byte 1
1 1 0 1 0 1 1 1 1 Byte 2
0 0 1 1 1 0 1 0 1 Byte 3
1 1 1 1 0 0 0 1 0 Byte 4
1 0 0 0 1 0 1 1 1 Byte 5
0 1 0 1 1 1 1 0 0 Check character (LRC)
⇑
Parity check (hardware generated)

Figure 6.16
The inclusion of a parity check with an LRC.

However, even when the two checks are used together, combinations of errors can still cause the error detection to fail. Figure 6.17 shows the effects of an even number of bit changes in the same columns of an even number of bytes. In this example, neither parity nor LRC errors are detected, even though the message has been received in error.

6.3.2
Cyclic Redundancy Check (CRC)

In certain applications, sending blocks of information with an LRC, VRC, or both, may provide sufficient error detection, especially if the channel is known to be reliable and reasonably error-free. However, some situations call for as near-to-perfect error detection as can be applied. Consider the electronic transfer of funds between bank branches. If, during a transfer, the channel is subject to noise, an undetected error can be disastrous for the bank and its customers since the amount received might not correspond to the amount sent. Although LRC and VRC can detect odd numbers of errors, it is possible for errors to go undetected. What is needed is an error detection technique that is several orders of magnitude more effective than LRC and VRC.

The technique used in many applications requiring better error detection is a cyclic redundancy check, or CRC.

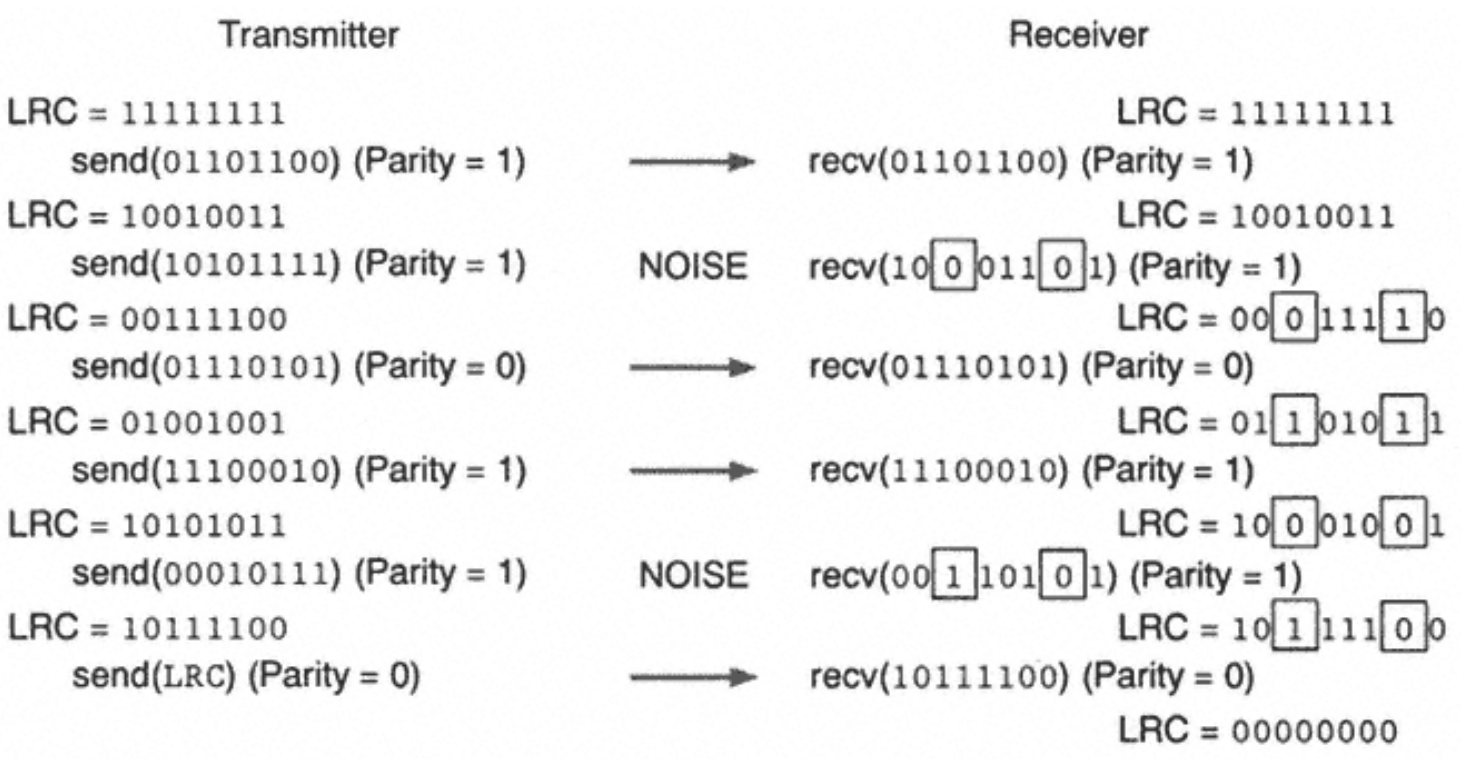

Figure 6.17
An error undetectable by both LRC and VRC.

CRC Fundamentals

The principles associated with the other error detection techniques discussed earlier are also applicable to the cyclic redundancy check algorithm. That is, the transmitting entity calculates the error check using the data from the message; appends a check sequence to the end of the message; and then transmits the message and the check sequence to the receiving entity in a frame. The receiving entity performs the same calculation on the incoming message and determines whether the message is valid by examining the result of the calculation. The cyclic redundancy check is particularly attractive because it offers greater error detection accuracy through the use of polynomial division.

Basically, the CRC algorithm treats the message as a single bit stream in which each bit is taken as a coefficient of a polynomial. In general, a message *k* bits long has *k* terms and is a polynomial of order *k*1. A message with terms *mk*1 + *mk*2 + · · ·+ *m*2 + *m*1 + *m*0 can be written as the polynomial:

$$M(x) = m_{k-1}x^{k-1} + m_{k-2}x^{k-2} + \cdots + m_2x^2 + m_1x^1 + m_0$$

For example, the message `101101001` can be written as the polynomial:

$$1x^8 + 0x^7 + 1x^6 + 1x^5 + 0x^4 + 1x^3 + 0x^2 + 0x^1 + 1x^0$$

and simplified to *x*8 + *x*6 + *x*5 + *x*3 + *x*0.

As with LRC and VRC, the message is transmitted with the check bits (commonly referred to as the CRC). Together, the message and the check bits make a unique cyclic code word. The objective of the CRC algorithm is to ensure that the vast majority of errors are detected by the receiving entity. By dividing the message polynomial by a generator polynomial, a set of check bits unique to the message can be obtained.

The generator polynomial, $G(x)$, is a polynomial of degree g. which must be less than the degree of the message polynomial, $M(x)$. G(x) is always odd (i.e., the lowest-order term has a value of 1) and must have a value greater than one. Unless otherwise indicated, all calculations are performed on binary digits using modulo 2 addition without carries ($\oplus$ denotes modulo 2 addition).

The transmission algorithm is as follows:

1. The message $M(x)$ is multiplied by xg (i.e., the message is shifted left by g bit positions; these bit positions are cleared).

2. The result of the multiplication is divided by $G(x)$, giving a quotient, $Q(x)$, and a remainder, $R(x)$:

$$\frac{x^g \times M(x)}{G(x)} = Q(x) \oplus \frac{R(x)}{G(x)}$$

$R(x)$ is always less than $G(x)$, since the maximum number of bits in the remainder is g.

3. $R(x)$ is added to the shifted message (i.e., the lower g bits), producing the frame to be transmitted, $T(x)$:

$$T(x) = x^g \times M(x) \oplus R(x)$$

The receiving entity receives the frame $T(x)$ and proceeds to divide the message by the generator polynomial, $G(x)$. Since the transmitted message includes the remainder from the original division, $T(x)$ should be exactly divisible by $G(x)$:

$$\frac{T(x)}{G(x)} = \frac{x^g \times M(x) \oplus R(x)}{G(x)}$$

Substituting $Q(x) \oplus \frac{R(x)}{G(x)}$ for $x^g \times \frac{M(x)}{G(x)}$, one finds

$$\frac{T(x)}{G(x)} = Q(x) \oplus \frac{R(x)}{G(x)} \oplus \frac{R(x)}{G(x)}$$

Since any number exclusive-or'ed with itself is 0 (such as $\frac{R(x)}{G(x)}$), a successful transmission will result in a remainder of 0.

Example

Consider the transmission of the message `101101001` using the generator polynomial `101001` ($x5 + x3 + 1$). Following the algorithm described above, one finds

1. $G(x)$ has $g = 5$, meaning that $M(x)$ must be shifted left by 5:

$2^5 \times$ `101101001` = `10110100100000`

2. The shifted message is then divided by the generator polynomial:

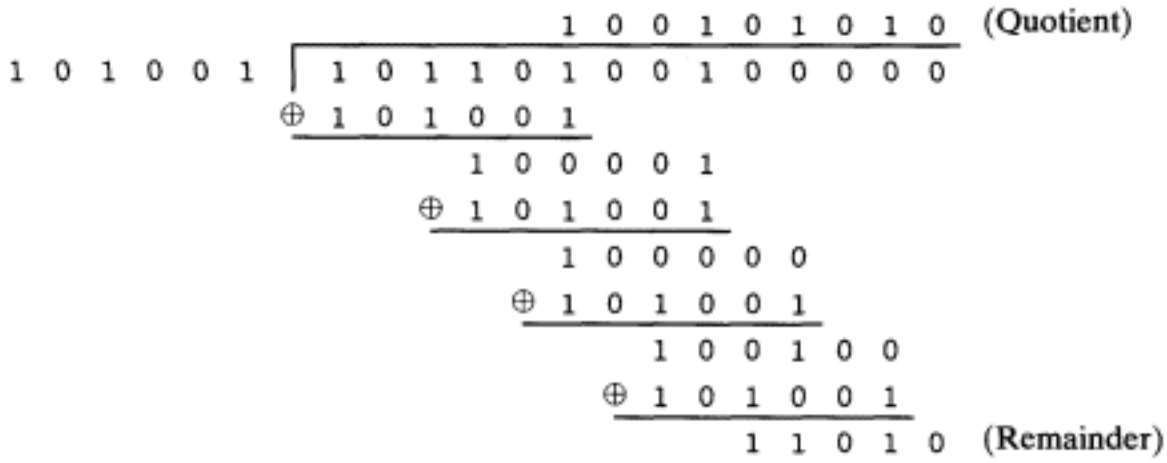

The quotient, $Q(x)$, is discarded.

3. $R(x)$, `11010`, is added to the shifted message, producing $T(x)$, consisting of $M(x)$ and $R(x)$. $T(x)$ is therefore `10110100111010`.

The receiving entity accepts the transmitted frame and performs the division using the same generator polynomial:

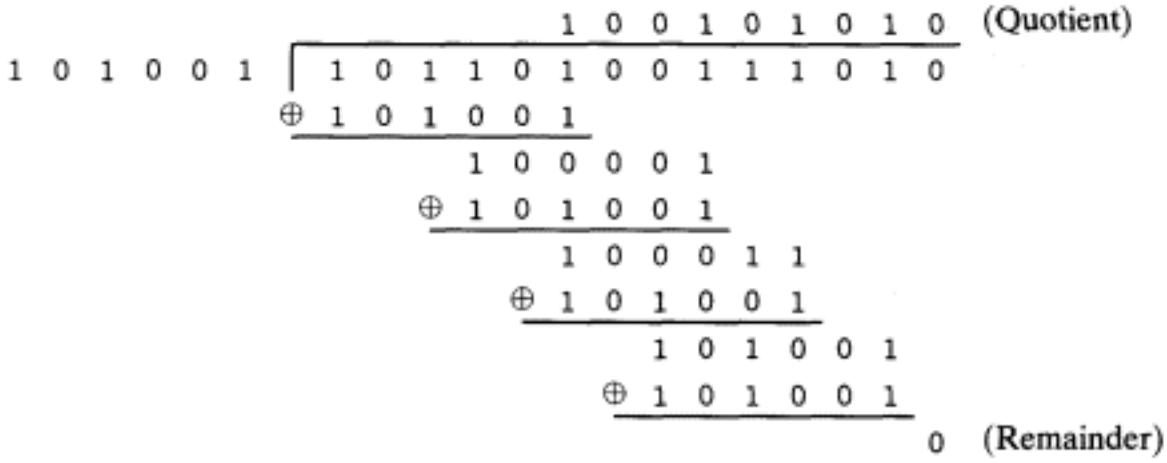

The result of the division is 0, indicating that no errors were detected in the frame.

However, should one or more bits become inverted (i.e., 0 to 1 or vice versa), the division should result in a non-zero remainder. For example, if the frame received was `101`**`01`**`100111010` (rather than `101`**`10`**`100111010`), the division would proceed as follows:

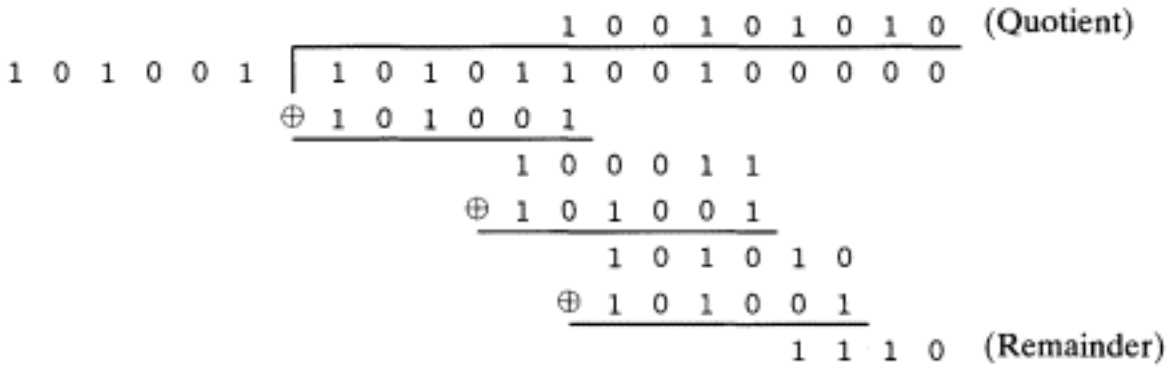

The receiving entity can reject the frame, since the remainder after the division is non-zero.

CRC Generators

All communication channels can be subject to some form of fault that leads to one or more bits in the transmitted frame becoming inverted. The CRC generator polynomial must be designed to detect invalid frames.

It is possible to represent the error itself as a polynomial that is exclusiveor'ed to the frame (if the bits in error are represented as coefficients, the corresponding bits in the frame will be inverted). Using the example from above in which the transmitted frame, `10110100111010`, is corrupted to `10101100111010`, the bits in error are as follows:

Transmitted frame `10110100111010`
Error bits `00011000000000`
Received frame `10101100111010`

In this example, the error polynomial $E(x)$ is $x10 + x9$. Error bits corresponding to polynomials containing $G(x)$ as a factor will be undetected. The question is how to develop a generator polynomial that will detect errors within the transmitted frame.

The CRC generator polynomial is chosen to detect the following four error types:

Single-bit errors. If the transmitted frame has a single bit in error, the error polynomial $E(x)$ contains a single term, xe. where e has a value less than the total number of bits in the frame $T(x)$. By having more than one term in $G(x)$, xe cannot be divided evenly, meaning that all single-bit errors will be detected.

Double-bit errors. A double-bit error, like the single-bit error, can be represented by the error polynomial. In this case, the polynomial consists of two terms, $E(x) = xi + xj$; as before, i and j are both less than the number of bits in the frame $T(x)$. $E(x)$ can be expressed as $E(x) = xi \times (xji + 1)$ if $i < j$. By ensuring that the generator polynomial has a factor with three

terms, neither *xi* nor (*xji* + 1) are divisible by the generator, meaning that all double-bit errors will be detected.

Odd number of errors. If the frame is received with an odd number of errors, $E(x)$ will have an odd number of terms. Since no polynomial with an odd number of terms has $x + 1$ as a factor, if $G(x)$ has $x + 1$ as a factor, all frames with an odd number of errors will be detected.

Error bursts. An error burst is a sequence of at least two incorrect bits separated by a series of other bits that may be in error. Error bursts less than or equal to *g* bits in length (i.e., the degree of the generator polynomial) will be detected as errors. If the burst is exactly $g + 1$ bits long (i.e., equal to the number of bits in the generator polynomial), the error will be detected if $E(x)$ does not equal $G(x)$. The probability of the error going undetected is $\frac{1}{2}^{(r-1)}$ (where *r* is the number of bits in the remainder), since the first and last bits of $E(x)$ are the same as $G(x)$, meaning that the remaining *r*1 bits must be equal. Finally, if the burst exceeds g, the probability of the error going undetected can be shown to be $\frac{1}{2}^{r}$.

The number of bits in the CRC is typically a function of the character (i.e., byte) size used by the communicating entities, since the supporting hardware handles quantities of this size. For example, in 8-bit communications, the remainder is typically 16 or 32 bits long, meaning that the generator polynomial can be 17 or 33 bits long.

There are a number of CRC generators available for use with synchronous communication protocols. Some of the better known are (note that the *X* in CRC-*X* refers to the number of bits transmitted in the remainder):

CRC-12a 12-bit CRC intended for use with synchronous protocols supporting 6-bit bytes. The generator polynomial is *x*12 + *x*11 + *x*3 + *x*2 + *x* + 1. The initial value of the CRC is 0. If, after the end-of-message is detected, the receiver's CRC result is 0, it is assumed that the message was received correctly.

CRC-16a 16-bit CRC used by IBM for the transmission of 8-bit character codes. The generator polynomial is *x*16 + *x*15 + *x*2 + 1. The initial value of the CRC is 0. The message is assumed to be correctly received if the receiver's CRC result is 0.

CRC-CCITTa 16-bit CRC developed by IBM for its SDLC protocol and also used by the ISO HDLC protocol (both protocols are discussed later); it employs the following generator polynomial: *x*16 + *x*12 + *x*5 + 1. The CRC is initialized to `0xFFFF`, and the CRC bits are inverted upon transmission. The received message is deemed correct if, after the end of the message, the receiver's CRC has a value of `0xF0B8`.

CRC-32a 32-bit CRC intended for use primarily in local area networks and some military applications. The generator polynomial is

$$x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$$

The initial value of the CRC is `0xFFFFFFFF`. As with the CRC-CCITT, the transmitter inverts the CRC bits prior to transmission. The receiving entity's CRC should equal `0xDEBB20E3` if the message is received correctly.

Although the generator polynomials are different, the basic algorithm is the same for each.

CRC Implementations

The first observation that one must make when considering the implementation of a CRC generation routine (in either hardware or software) is *how* is it possible to perform division on a message that may be thousands of bits long, when most machines can perform only 32-bit arithmetic (or less in many cases).

Interestingly enough, the entire division process can take place in a shift register and can be performed on each bit as it is being transmitted serially. Figure 6.18 shows a possible shift register design for CRC-16 calculations (polynomial generator: *x*16 + *x*l5 + *x*2 + 1).

Before the first byte is transmitted, the CRC shift register is initialized to `0x0000`; it contains the intermediate remainder. The shift register emulates polynomial division by performing modulo-2 subtraction on the terms *x*0, *x*2, and *x*15. The result of the exclusive-or between the low-order data bit and the high-order bit of the remainder (*x*16) is either a 0 or a 1: a 0 causes the remainder to rotate unchanged, whereas a 1 inverts the value of each divisor polynomial termequivalent to obtaining a new intermediate remainder value by including bits from the dividend. Each bit from the Data Shift Register is put onto the channel for transmission as well as being fed into the CRC circuit. When the last data bit has been included in the intermediate remainder, the bits in the shift register are clocked onto the channel.

The Commkit CRC Utilities

Commkit is supplied with two CRC utilities that generate CRC-16 remainders. The first is a software emulation of the CRC shift register circuit described above, while the second performs a table lookup:

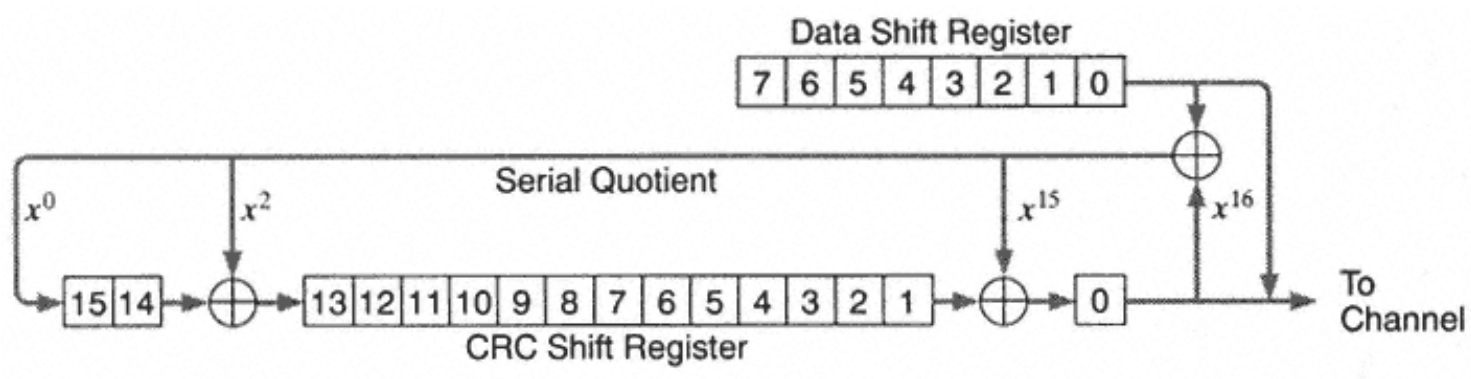

Figure 6.18
A shift-register implementation of CRC-16.

Emulation. The shift register shown in Figure 6.18 can be emulated in software by looping through each bit in the data byte prior to the transmission of the data.

`Crc_calc()` (found in the file `crcsr.c`) contains an implementation of the shift register written in C. The software is optimized upon entry to `crc_calc()`; an intermediate value of the remainder is produced by exclusive-or'ing of the data `byte` with `crc`. The low-order bit in the new intermediate remainder is right-shifted. Whenever a carry is detected, the polynomial `POLY` is exclusive-or'ed into the remainder (note that `POLY` is simply the octal representation of *x*15 + *x*2 + x0).

The intermediate remainder is returned by `crc_calc()`.

Table lookup. By studying the emulation software, one can see a number of patterns emerge that can reduce the time associated with division by repeated shifting. These patterns allow the creation of a table that can produce the next intermediate remainder from the current remainder and the data byte. A second, faster version of `crc_calc()` using table lookup can be found in `crctbl.c`.

The table is based upon the observation that each bit in the new intermediate remainder is always made of the same combinations of data bits and the previous value of the remainder. For example, bit 15 of the new remainder always contains the data bits exclusive-or'ed with the low-order bits of the old intermediate remainder. Similarly, bits 1 through 5 of the new remainder are the shifted values of bits 9 through 13 of the old remainder. The table `crc_tbl` is an array of 256 integers, each of which represents the different intermediate remainder results associated with the subscript. The table is created in `gen_tbl()`.

Upon entry into `crc_calc()`, the byte to be included in the CRC is exclusive-or'ed with the lower 8 bits of the CRC, producing an index into the table. The CRC is then shifted right and exclusive-or'ed with the bit pattern from `crc_tbl`; this is the new intermediate CRC.

Prior to the first call to `crc_calc()`, `gen_tbl()` must be called to generate the array `crc_tbl`.

Not surprisingly, both methods produce the same results for the same messages; however, table lookup is the faster of the two methods.

6.4 Error Recovery

The previous section showed a number of techniques that allow a receiving DTE to determine whether a frame has been received in error. The question now is: once the receiving DTE has detected the error, what steps should be taken to correct it?

Since there is only one transmitter and one receiver, the possible places where the error can be corrected are somewhat limited:

The receiving DTE can request the transmitting DTE to retransmit the frame in error; this is known as backward error recovery.

The receiving DTE can attempt to correct the error itself using what is known as forward error recovery.

6.4.1
Backward Error Recovery

Because of the distances and speeds involved, most terrestrial communications do not expect the receiving DTE to correct errors detected within the frame. Instead, most protocols are written so that the receiving DTE *discards* the frames in error and requests a retransmission of the frame using a technique known as automatic repeat request (ARQ). There are a number of ARQ algorithms used in backward error recovery, as illustrated by the following protocols.

Stop-and-Wait

A stop-and-wait protocol requires the transmitting entity to send one frame at a time and await a reply from the receiving entity. The reply is an acknowledgment, indicating whether the frame has been received correctly.

If the frame is correctly received, the receiving entity responds with a positive acknowledgment (or *ACK*). Figure 6.19 is a time sequence diagram that illustrates the basic cycle of a stop-and-wait protocol (unless otherwise indicated, it is assumed that the frame the DTE is about to transmit contains a message).

In Figure 6.19, DTE 1 gets a message, transmits it in a frame, and then waits for a reply from DTE 2. When DTE 2 receives the error-free message, an acknowledgment is returned. Upon receipt of the *ACK*. the transmitting entity (DTE 1) can discard the message just sent and get the next message for transmission.

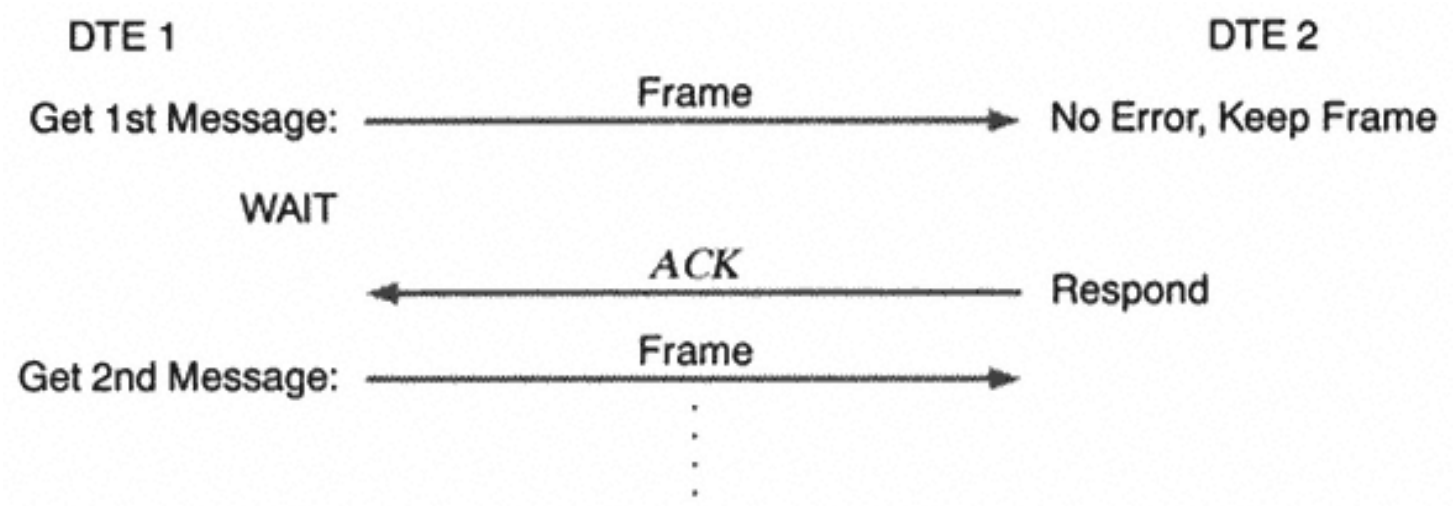

Figure 6.19
A time-sequence diagram for the stop-and-wait protocol.

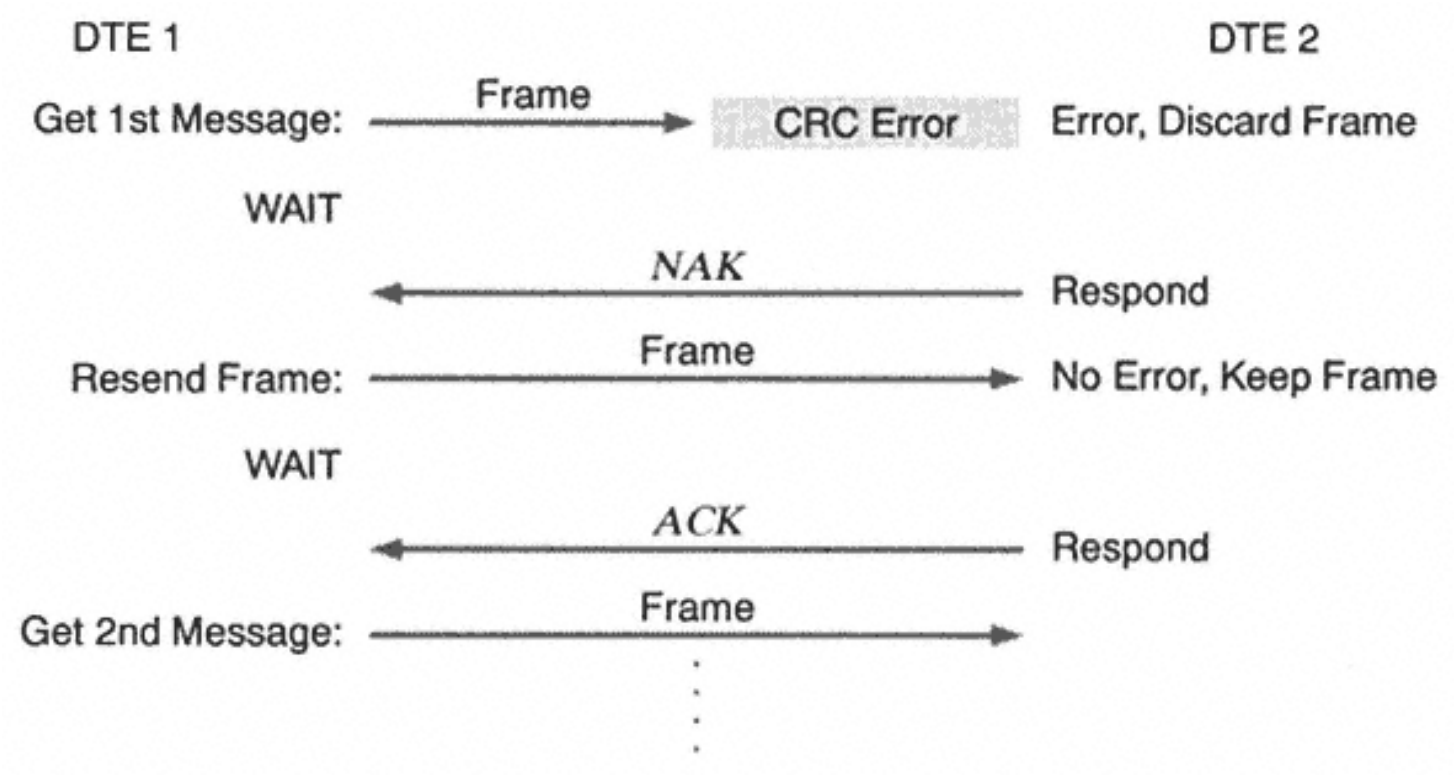

Figure 6.20
Error detection and correction in the stop-and-wait protocol.

However, if a frame is received incorrectly, the receiving entity (DTE 2) responds with a negative acknowledgment (or *NAK*), requiring the transmitting entity to retransmit the *original* frame. As shown in Figure 6.20, once the frame is received correctly, transmission can continue with the subsequent frames.

The transmit-frame-receiveacknowledgment (positive or negative) cycle continues until all messages have been sent (in frames) and acknowledged.

The algorithms for the stop-and-wait protocol can be written as follows:

Transmitting data:

1. Send a frame (containing a message) to the remote entity.
2. Wait for an acknowledgment.
3. If the acknowledgment is positive, get the next message. Otherwise, keep the same message for retransmission.
4. Repeat steps 1, 2, and 3 while there are more messages to send.

Receiving data:

1. Wait for a frame.
2. If no errors are detected, keep the frame (extract the message) and respond with a positive acknowledgment.
3. If errors are detected, discard the frame and respond with a negative acknowledgment.
4. Repeat steps 1, 2, and 3 while there is more to be received.

The stop-and-wait protocol as described assumes that *all* frames and acknowledgments reach their intended destination. This might be an incorrect assumption, since the flow of data may cease if a frame or an acknowledgment (positive or negative) is lost (because of noise or a momentary loss of connection), leaving the transmitting entity waiting *indefinitely* for an

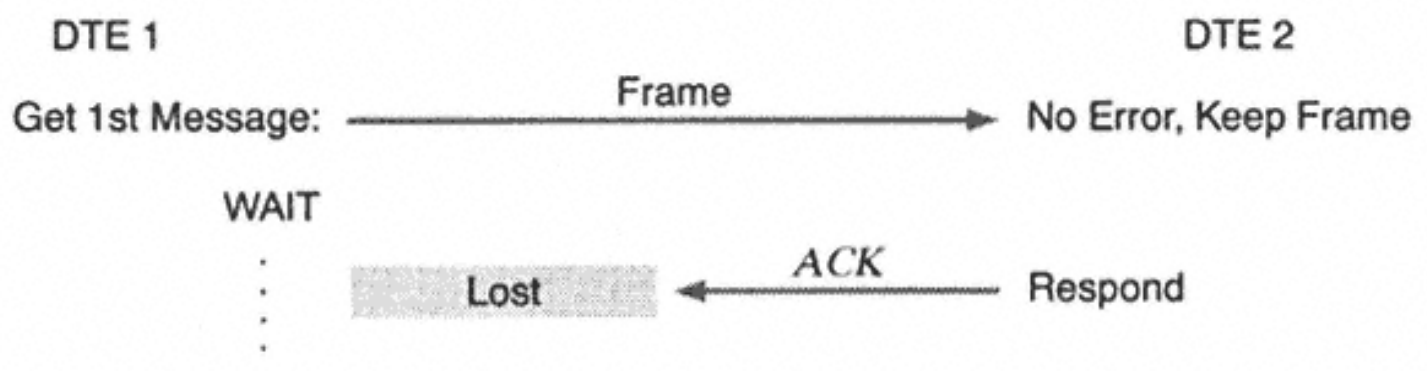

Figure 6.21
Information flow ceases if *ACK* is lost.

acknowledgment. An example of the loss of a positive acknowledgment is shown in Figure 6.21.

There are a number of ways in which the absence of an acknowledgment can be overcome to ensure the continuing flow of frames. The most common is to limit the amount of *time* the transmitting entity can wait for a response. That is, if a response isn't received within a specified time period, the transmitting entity simply retransmits the frame, as shown in Figure 6.22.

This solution, however, can lead to another problem: DTE 2 has now received a *second* copy of DTE 1's original frame (meaning that the same message has been received twice). The question arises: how can DTE 2 distinguish between frames that are retransmissions of previous ones and frames that are entirely new?

One possible solution is to have DTE 2 compare the contents of the previous frame (i.e., the message) with the one just received: if they are the same, the one just received can be discarded. This is not necessarily a good idea, since it ignores the possibility that two or more messages could have the same value, such as a pair of blank lines in a file. An alternative is to have each frame associated with an identifier that allows the receiving DTE to distinguish between incoming frames. For example, the identifier could be an integer that is sent with the message as part of the frame and is unique to each frame (see Figure 6.23). The identifier is more commonly referred to as a sequence number.

Duplicate frames can be identified and rejected by, for example, the receiving DTE examining the sequence number associated with each frame, as in Figure 6.24. In the figure, DTE 1 sends frame #0 and waits for a reply.

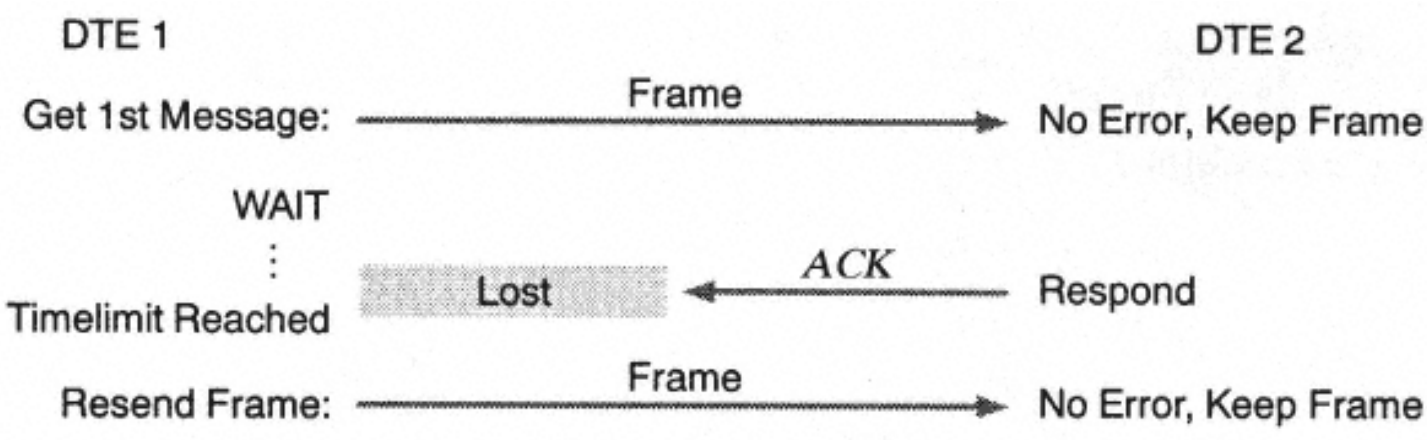

Figure 6.22
Using a time limit to detect a missing message.

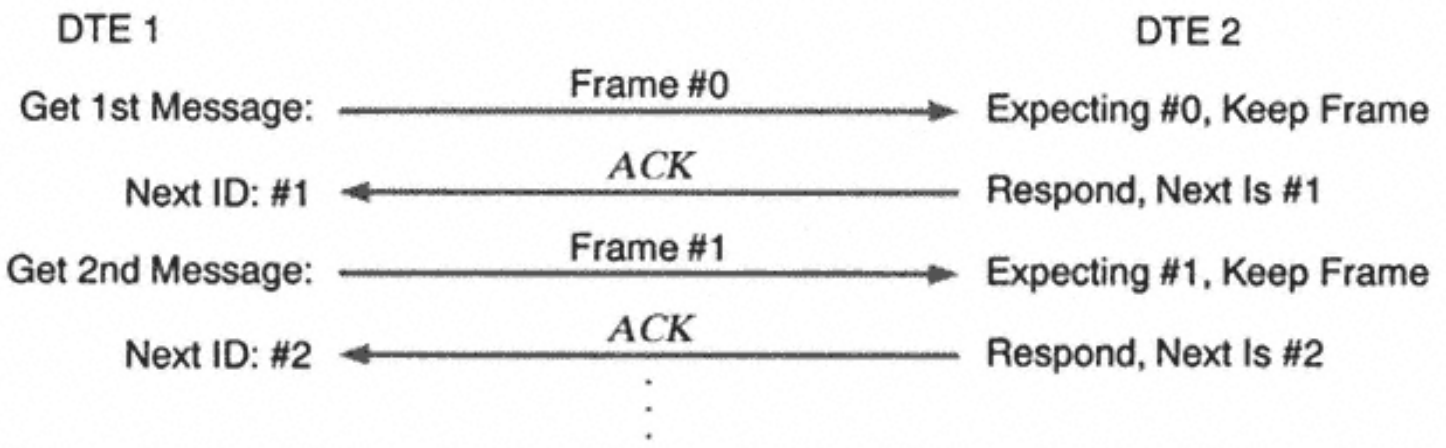

Figure 6.23
Associating each frame with an identifier.

Since no reply is received before the time limit is reached, frame #0 is retransmitted. DTE 2 initially receives frame #0 and responds with an *ACK*. since 0 is the expected sequence number; the expected sequence number is then increased to 1. When the next frame arrives with a sequence number less than the expected sequence number, DTE 2 correctly assumes that the frame is a duplicate and discards it. The frame is a duplicate because the previous *ACK* was lost, meaning that DTE 2 must retransmit the acknowledgment.

Now, when a duplicate frame is received, the receiving DTE need only inspect the sequence number to determine whether the frame is new or a retransmission of the previous one. As before, when the transmitting DTE finally receives a positive acknowledgment, the next frame can be transmitted.

It may appear that each frame sent must have its own unique sequence number (for example, starting at 0 and proceeding to infinity), but careful examination of the transmissionacknowledgment cycle reveals that, at any particular moment, only a single frame must be acknowledged. This means that the stop-and-wait protocol needs at most *two* distinct sequence numbers: one for the frame about to be sent and the other for the previous frame. By alternating between two sequence numbers (say, 0 and 1), the two DTEs can remain in step and determine which frames are lost or received in error.

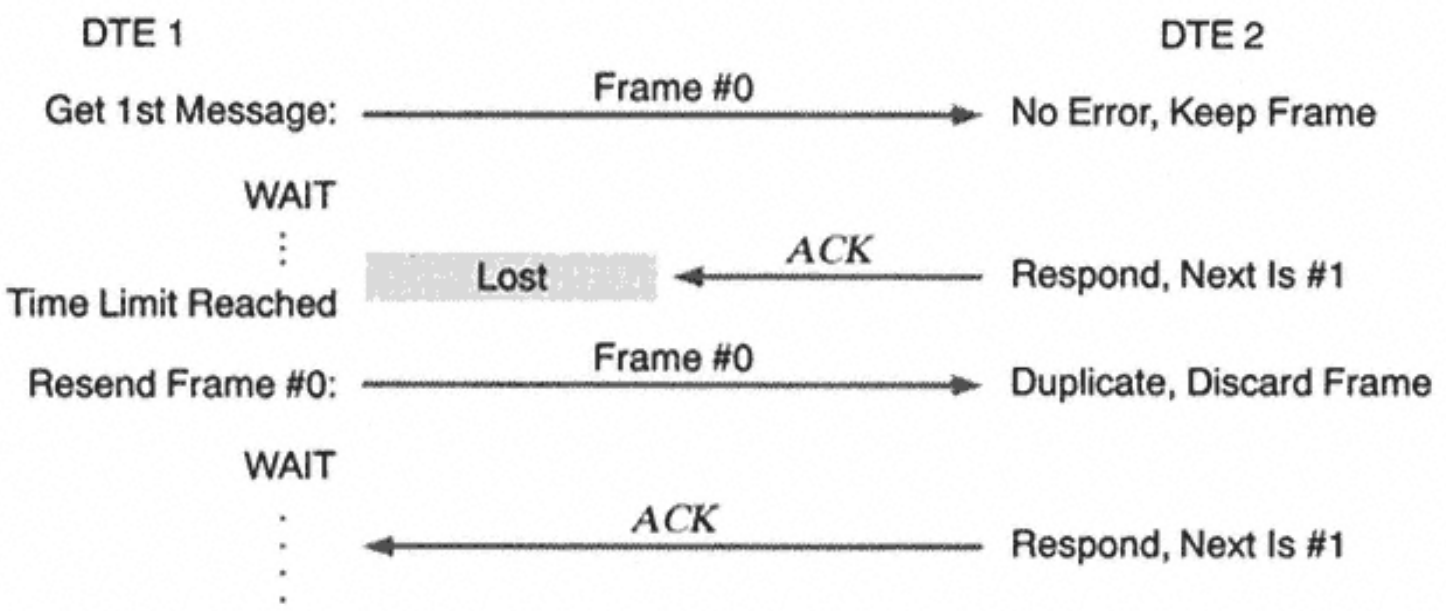

Figure 6.24
Recognizing a duplicate frame by its sequence number.

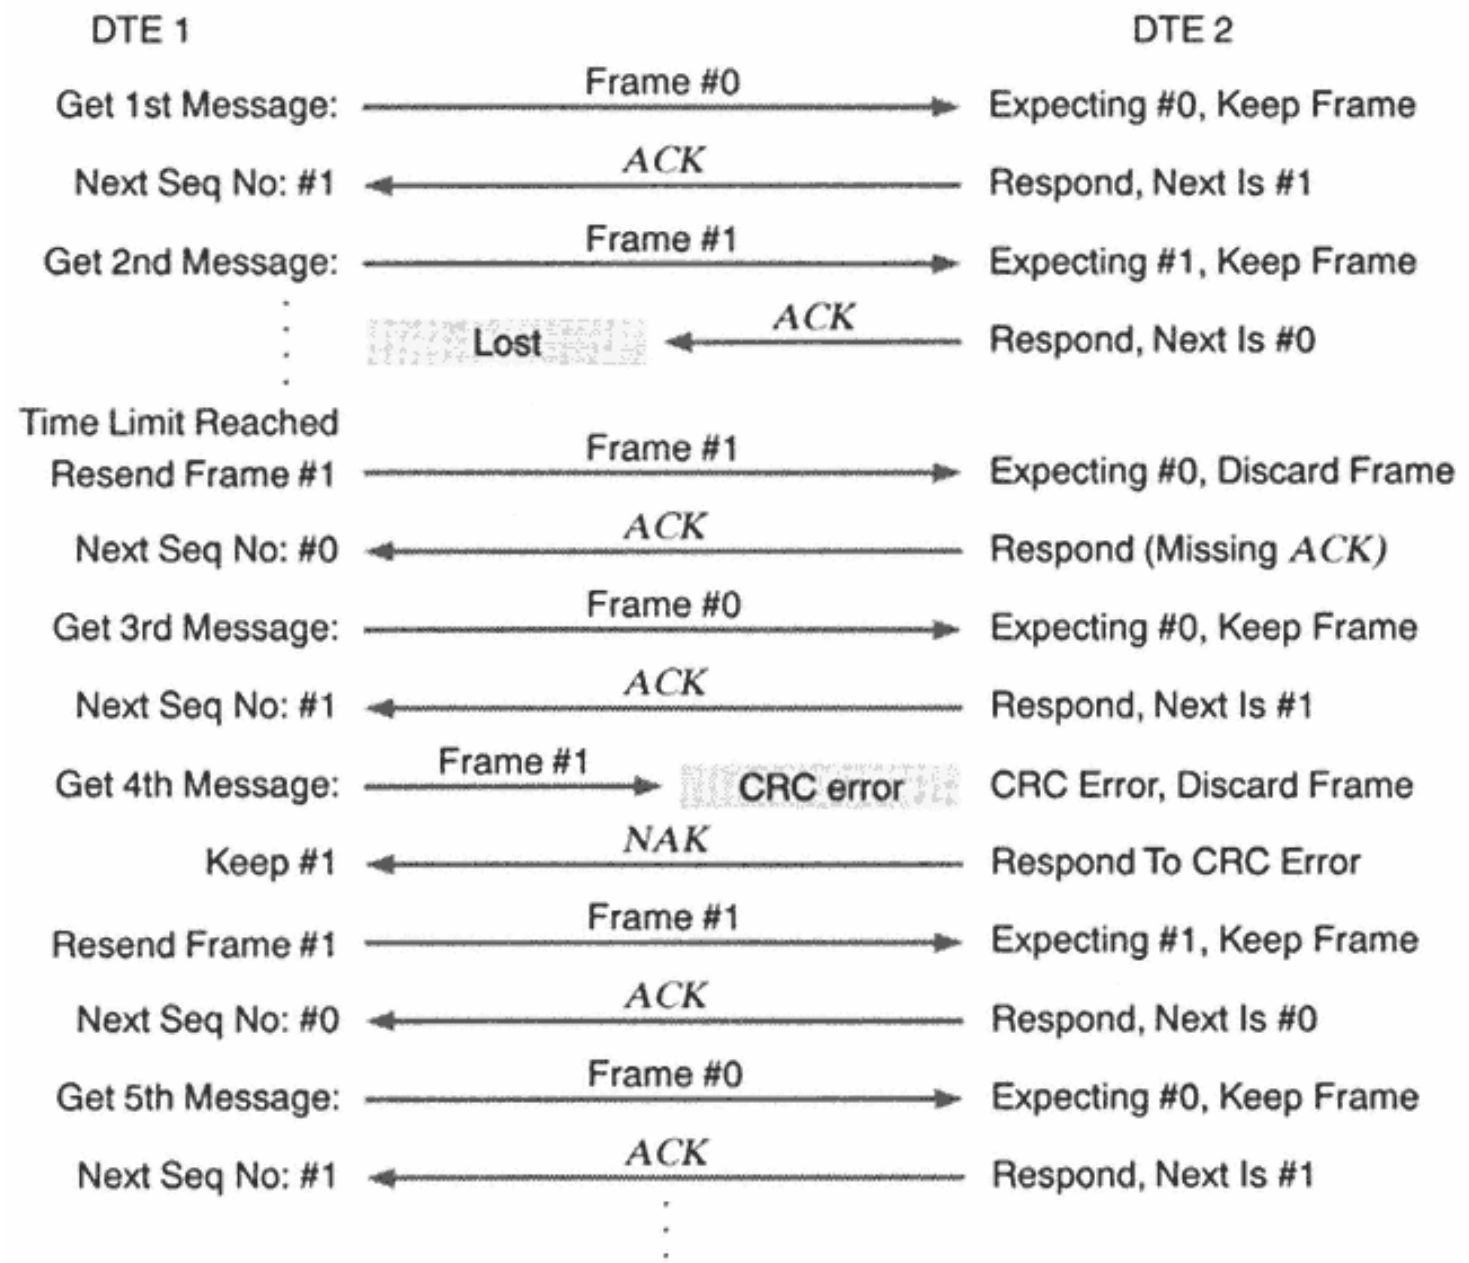

Figure 6.25
An example of the stop-and-wait protocol.

Figure 6.25 is an example of the stop-and-wait protocol transmitting five messages using two sequence numbers.

The receiving DTE must maintain state information regarding the frame it is *about* to receive (that is, the expected sequence number). Furthermore, both DTEs must agree to a common starting sequence number (typically 0). A complete stop-and-wait transmission-state machine is shown in Figure 6.26 (the first frame is sent with sequence number 0 and enters state T2).

In the diagram shown in Figure 6.26, state T2 is entered after frame #0 has been sent (either initially or after receiving an *ACK* in state T1). The

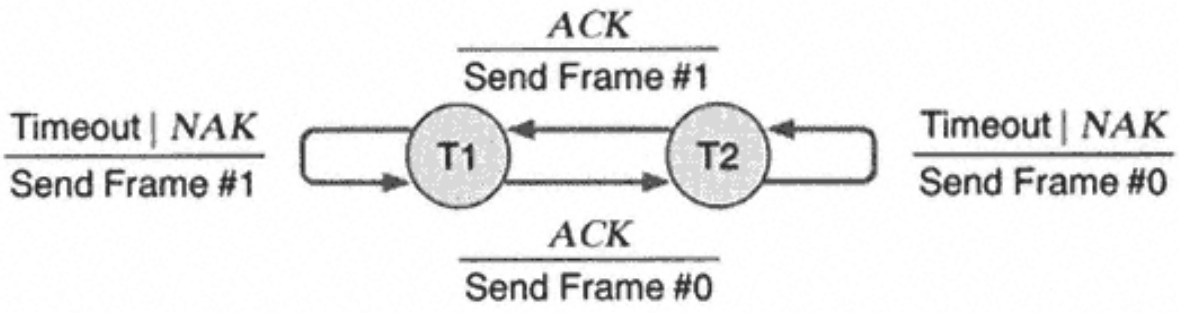

Figure 6.26
Stop-and-wait transmission state diagram.

machine stays in state T2 if a timeout occurs or a *NAK* is received, either of which events cause frame #0 to be retransmitted. Upon receipt of an *ACK* in state T2, frame #1 is sent and control passes to state T1. A timeout or a *NAK* causes frame #1 to be retransmitted from state T1, while an *ACK* results in a state change (to T2) and the transmission of frame #0. Although not shown, the machine should terminate upon receipt of an acknowledgment of the last frame.

A stop-and-wait receive-state machine for the stop-and-wait protocol is shown in Figure 6.27. It consists of two states: R1 (to be entered initially while waiting for frame #0 or after receiving frame #1), and R2 (entered after receiving frame #0). In both states, the reception of a frame with a bad CRC causes the transmission of a negative acknowledgment (*NAK*), while the receipt of a valid frame causes transmission of a positive acknowledgment (*ACK*). Receiving a valid frame after it has already been acknowledged (for example, receiving frame #1 in state R1) indicates that the transmitter did not receive the previous acknowledgment, requiring the transmission of an *ACK*. The reception of a bad CRC in either state could be a retransmission of a previous frame (if the last *ACK* was not received), or it could be the transmission of a new one. For example, in state R2, if frame #0 is received with an invalid CRC, the last *ACK* was not received and the retransmission was in error. If frame #1 is received with a bad CRC in state R2, the last frame was correctly acknowledged, but frame #1 was damaged during its transmission. In either event, the response is a *NAK*.

Interestingly enough, the stop-and-wait protocol does not require the negative acknowledgment: if the DTE receives a frame with a CRC error, there is no need to respond since the transmitter will eventually time out and retransmit the same frame again. The advantage of using the negative acknowledgment is that throughput can be increased because a frame with a bad CRC is retransmitted as soon as the transmitting DTE receives the *NAK*.

Go-Back-N

A closer examination of the stop-and-wait protocol reveals that it is essentially half-duplex. For example, after the transmitting entity sends a frame, the channel is turned around to allow the receiving entity to respond with an

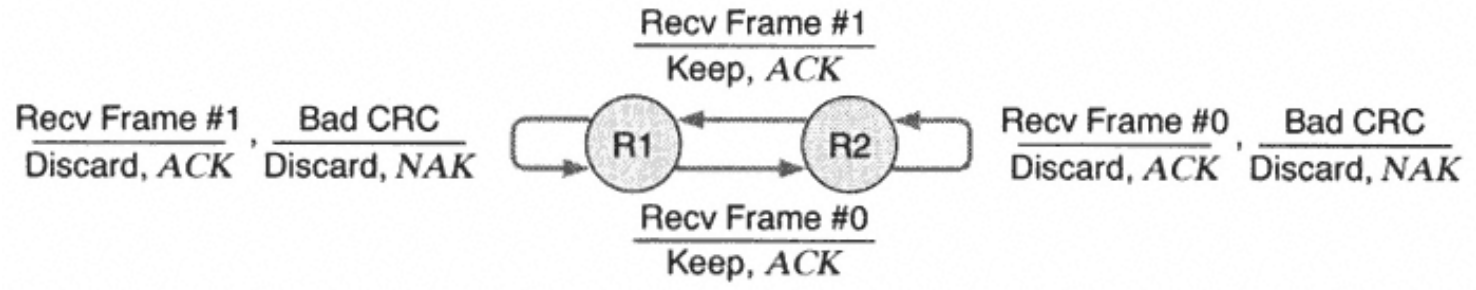

Figure 6.27
Stop-and-wait receive-state diagram.

acknowledgment. If the connection between the two entities is full-duplex, using the channel in a half-duplex fashion is a waste of time and bandwidth. The question now iswhether a full-duplex communication channel can be used to increase throughput?

One way of increasing throughput is to allow the transmitting DTE to send a number of frames while simultaneously receiving responses. Figure 6.28 is an example of such a scenario (the channel is assumed to be full-duplex, since frames and acknowledgments are transmitted simultaneously; the crossover of a frame and an acknowledgment indicates simultaneous transmission and reception).

In Figure 6.28, DTE 1 discards frame #0 upon receipt of the first acknowledgment and frame #1 upon receipt of the second acknowledgment. This approach will work correctly as long as no frames are lost, as Figure 6.29 illustrates. There, DTE 1 receives an acknowledgment and discards frame #0something that shouldn't happen because the first frame has been lost, never reaching DTE 2; the acknowledgment that DTE 1 receives is for frame #1. What is needed is a mechanism by which the acknowledgment can indicate the specific frame being acknowledged rather than a single acknowledgment being used to acknowledge all frames.

A commonly adopted solution is to append an identifier to the acknowledgment (in much the same way as the identifier is appended to the frame). The identifier indicates the sequence number of the *next* frame that the receiving DTE expects to receive. This is analogous to sending letters to someone, then telephoning and asking which letters they have receivedthe person might reply that the first and second letters have been received, but not the third; or, if no letters have been received, the response might be that the first one has yet to arrive. In the case of the transmitting DTEs, a

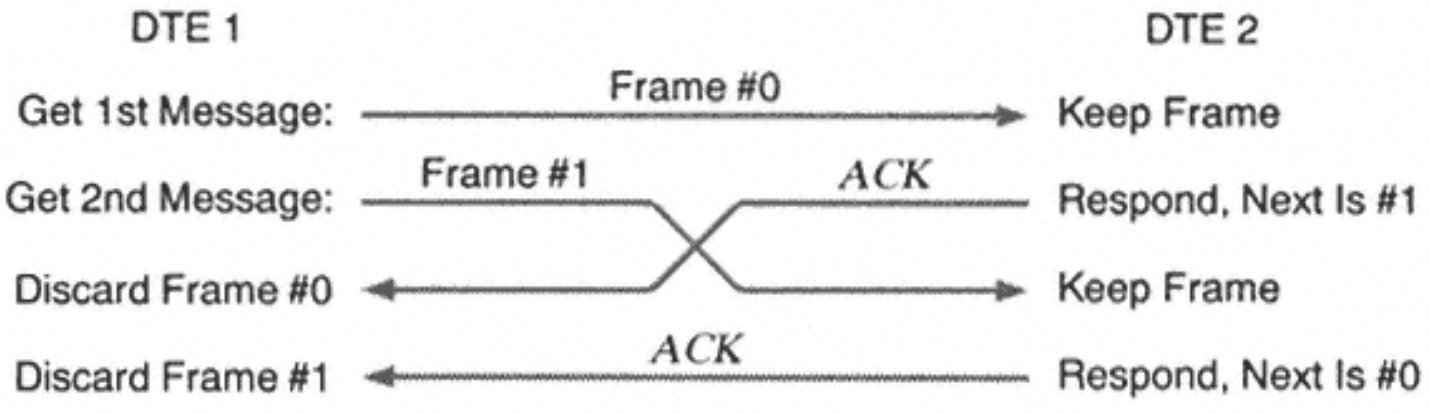

Figure 6.28
The simultaneous transmission of frames and acknowledgments.

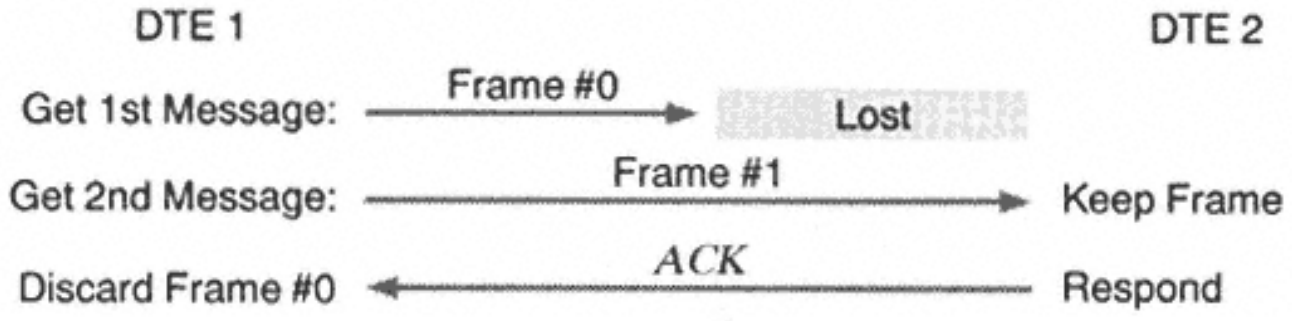

Figure 6.29
The transmitter discarding the wrong frame.

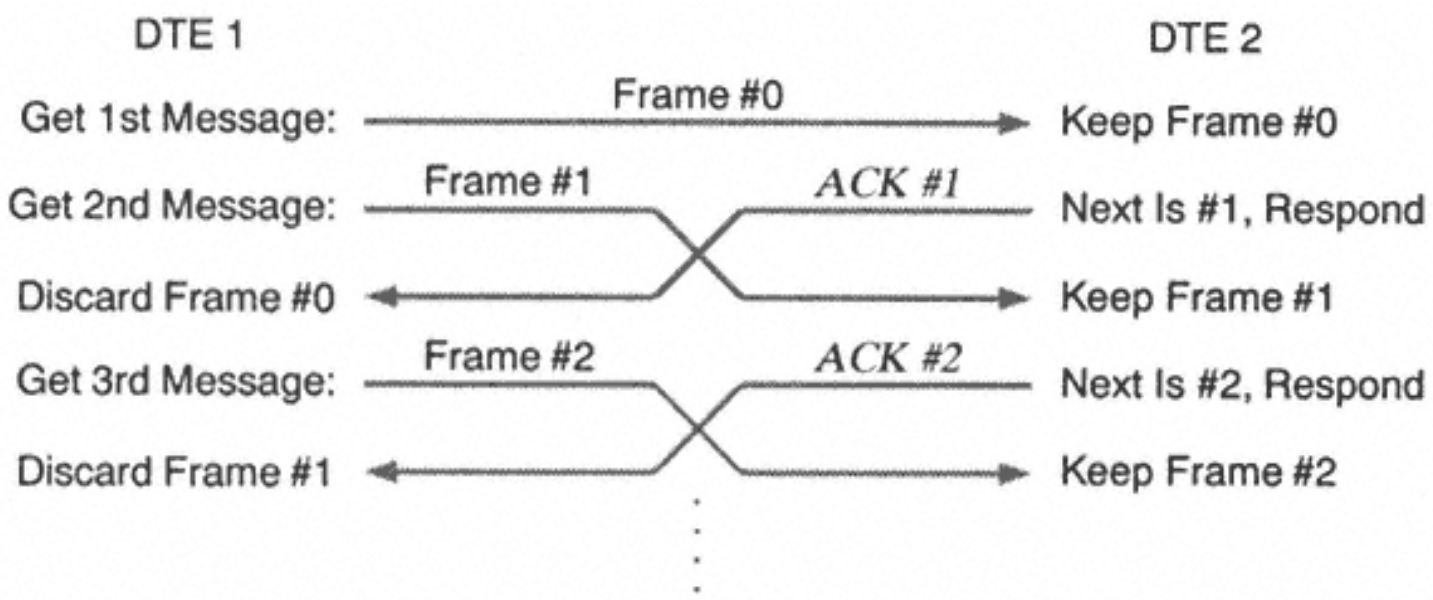

Figure 6.30
Adding a sequence number to the acknowledgment.

typical transmission scenario can be as shown in Figure 6.30 (note that the acknowledgments indicate the sequence number of the *next* expected frame).

Upon receipt of an acknowledgment, the transmitting DTE can discard the frames with sequence numbers up to (but not including) the sequence number sent with the acknowledgment, since the acknowledgment is indicating the *next* frame expected by the receiver. For example, if an acknowledgment with sequence number 3 is received, frames #0, #1, and #2 can be discarded (if this hasn't already been done).

Both the transmitting and receiving DTEs must maintain sufficient information to determine the current state of the communication (i.e., the number of the next frame to transmit, the number of the next expected frame, and so on). This information is represented as a pair of state variables:

Vs. The send state variable (sometimes written as V(S)) is the sequence number of the next frame to be transmitted; the initial value of Vs is typically 0. Vs is incremented as each new frame is transmitted.

Vr. The receive state variable (also written as V(R)) is the sequence number that is next expected by the receiver; Vr has an initial value identical to that of Vs (i.e., 0). Receipt of a frame with a frame number equal to Vr causes Vr to be incremented.

The sequence number associated with each frame is referred to as the send sequence number and is written as Ns (or N(S)). The transmitter assigns the value of Vs to Ns; once assigned, the value of Vs is incremented.

The sequence number appended to the acknowledgment, the receive sequence number, Nr (or N(R)), is the value of Vr after the frame has been accepted by the receiver. To the receiving DTE, Nr is the value of the next expected sequence number, while to the transmitting DTE all frames with sequence numbers less than Nr have been accepted by the receiving DTE.

As an example, Figure 6.31 shows the same sequence of events that took place in Figure 6.30, this time using the state variables Vs and Vr and the sequence numbers Ns and Nr. Prior to the start of a communication, both the

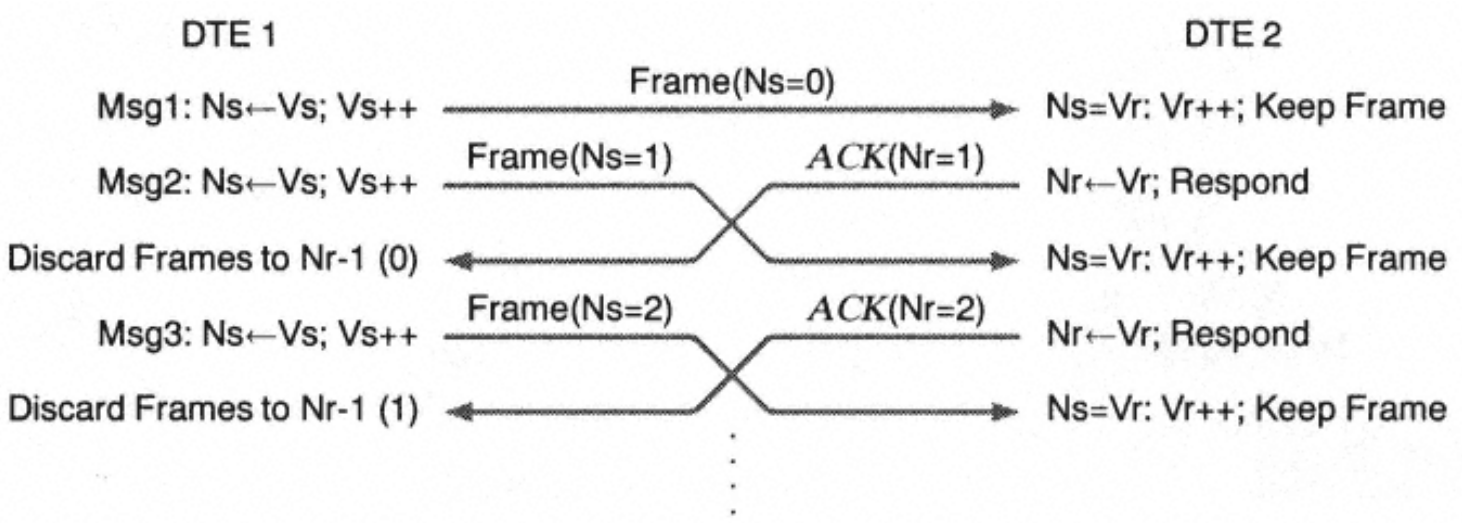

Figure 6.31
Using the state variables (Vs and Vr) and sequence numbers (Ns and Nr).

transmitting and receiving DTEs must ensure that their values of Vs and Vr are identical (in this and subsequent examples, both are assigned the value 0). DTE 1 sends the first frame, which contains message 1, and Ns, which has been assigned the current value of Vs, or 0 (the assignment written as Ns ← Vs). Vs is incremented by 1 (this is denoted by Vs++); DTE 1 keeps the frame until its reception is acknowledged. Upon receipt of the first frame, DTE 2 compares the value of Vr with the value of Ns taken from the frame; since both are equal (indicated by "Ns = Vr"), the value of Vr is incremented (from 0 to 1) and the contents of the frame are kept.

DTE 2 responds with an acknowledgment; the value of Nr is obtained from the current value of Vr (in this case, 1). When DTE 1 receives the acknowledgment, all frames with sequence numbers less than Nr are discarded (i.e., frame "Ns = 0"). This cycle continues until DTE 1 has no more messages to transmit.

Including sequence numbers with acknowledgments ensures that the transmitting DTE "knows" which pending frames to discard. However, there are three other possible error conditions that must be considered: the loss of an acknowledgment, the loss of a frame, and the lack of a response from the receiving DTE.

If an acknowledgment is lost, the transmitting DTE continues to transmit (since there is no indication of the loss) and when the next acknowledgment is received, the transmitting DTE can determine which frames are to be discarded. This problem is illustrated in Figure 6.32 (both Vs and Vr are assigned 0 initially).

In Figure 6.32, the acknowledgment for frame (Nr = 0) (i.e., "*ACK* (Nr = 1)") is lost, but DTE 1 continues to transmit frames. Upon receiving "*ACK*(Ns = 2)," DTE 1 can dispose of both frame Ns = 0 and Ns = 1, since the acknowledgment indicates that *all* frames with sequence numbers less than Nr = 2 have been accepted.

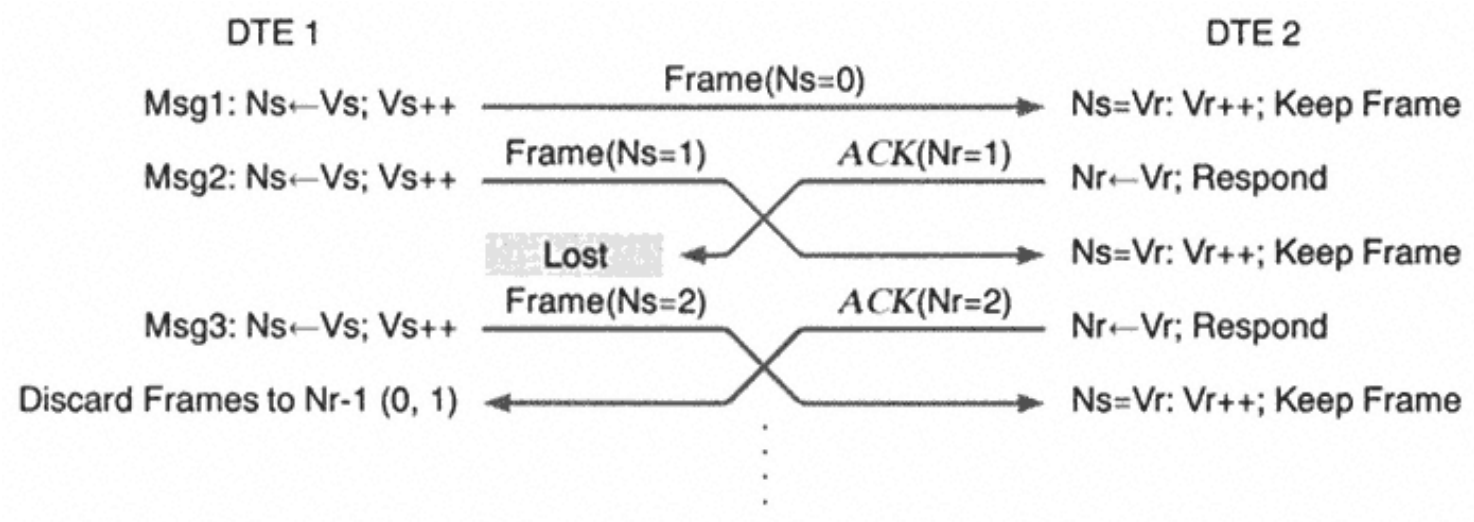

Figure 6.32
The receive sequence number (Ns) includes all previous sequence numbers.

The second error condition that must be considered is how to handle the loss of a frame. If a DTE receives a frame with an invalid or out-of-sequence value of Ns (identifiable because Ns does not equal Vr), the receiving DTE should respond with an indication that the frame's sequence number is invalid, thereby permitting the transmitting DTE to restart its transmission with the correct frame. The transmitting DTE is to resume with the frame that the receiving DTE next expects. By responding with a negative acknowledgment (and the sequence number of the next expected frame, that is, Vr), the transmitting DTE can distinguish between a normal acknowledgment and an error condition. In Figure 6.33, frame (Nr = 0) is lost, causing DTE 2 to respond with a request for frame (Nr = 0).

The negative acknowledgment received by DTE 1 indicates that DTE 2 is still expecting frame (Ns = 0), even though frames Ns = 0, Ns = 1, and Ns = 2 have been sent. Since DTE 1 has no indication as to which frames have been received correctly and which are in error, the worst is assumed and transmission resumes from the specified frame (in this case, frame (Ns = 0)). Note that during the retransmission, the value of Ns in each packet remains unchanged and Vs is not incremented.

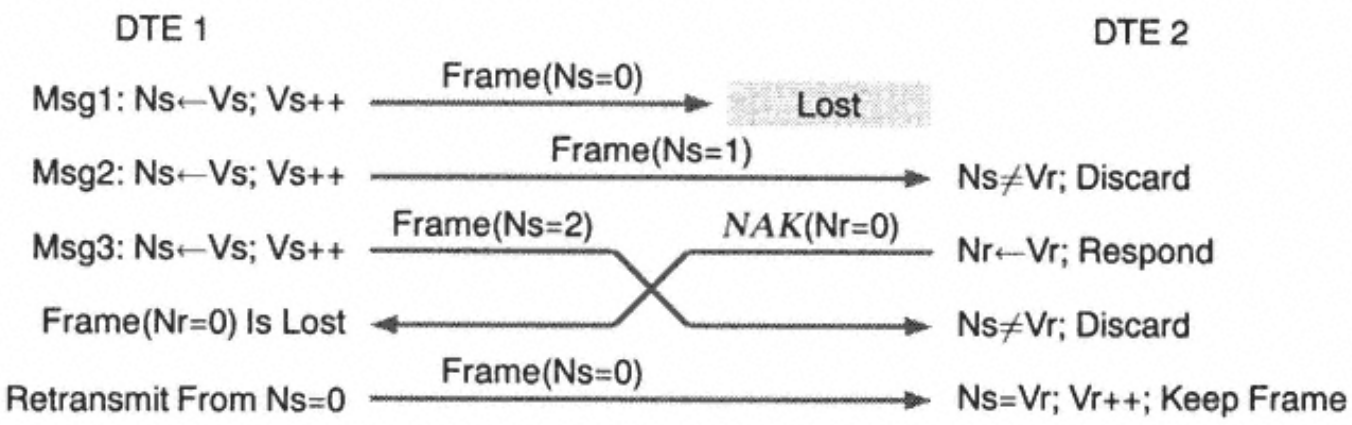

Figure 6.33
The retransmission of a lost frame.

Protocols that expect the transmitting DTE to retransmit all frames from the sequence number indicated by the receiving DTE are known as go-back-N. In a go-back-N protocol, all frames that are received with an invalid or unexpected value of Ns are discarded.

The third error condition occurs when the transmitting DTE sends a frame and does not receive a response. The solution adopted in the stop-and-wait protocol can also be applied to the go-back-N protocol: each frame that is sent must be acknowledged within a certain time period; otherwise, the frame is resent. Since several frames may be awaiting acknowledgment, each one is associated with a unique timer; as acknowledgments are received, the timers are turned off and the frames discarded.

There are two different scenarios where frame loss may occur. In the first, the frame does not arrive at the receiving DTE; this is illustrated in Figure 6.34. When the timer associated with frame (Ns = 0) expires, DTE 1 retransmits the frame, including its original value of Ns; Vs remains unchanged, since this is a retransmission.

The second scenario occurs when the frame has arrived and is acknowledged by the receiving DTE, but the acknowledgment is not received by the transmitter. In this situation, when the time limit expires, the frame is retransmitted but discarded by the receiving DTE. As shown in Figure 6.35, DTE 2 responds with an acknowledgment indicating the next expected sequence number. Note that the transmitting DTE (i.e., DTE 1) cannot distinguish between the two scenarios.

To avoid unnecessary retransmissions, the timer associated with the acknowledged frame is turned off as soon as the correct acknowledgment is received. Should several frames be waiting for acknowledgments, and the incoming acknowledgment indicates that all of the frames have been accepted, each timer is turned off, and the associated frame is discarded.

One of the arguments for the go-back-N protocol is that it can be implemented on a full-duplex channel, thereby increasing channel throughput. However, two competing objectives must be considered in this case

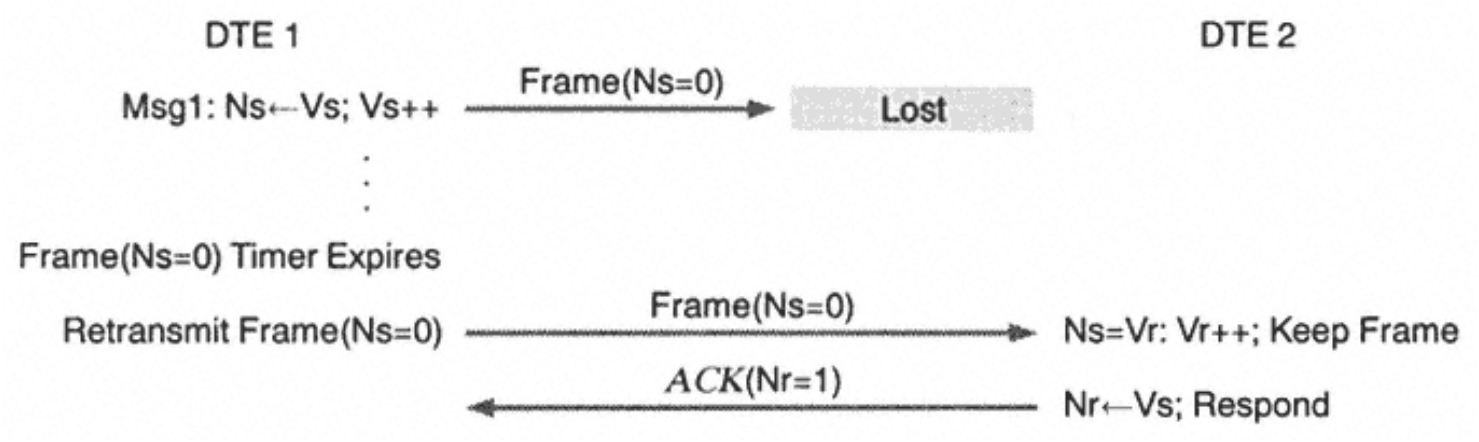

Figure 6.34
Using a timeout to retransmit a lost frame.

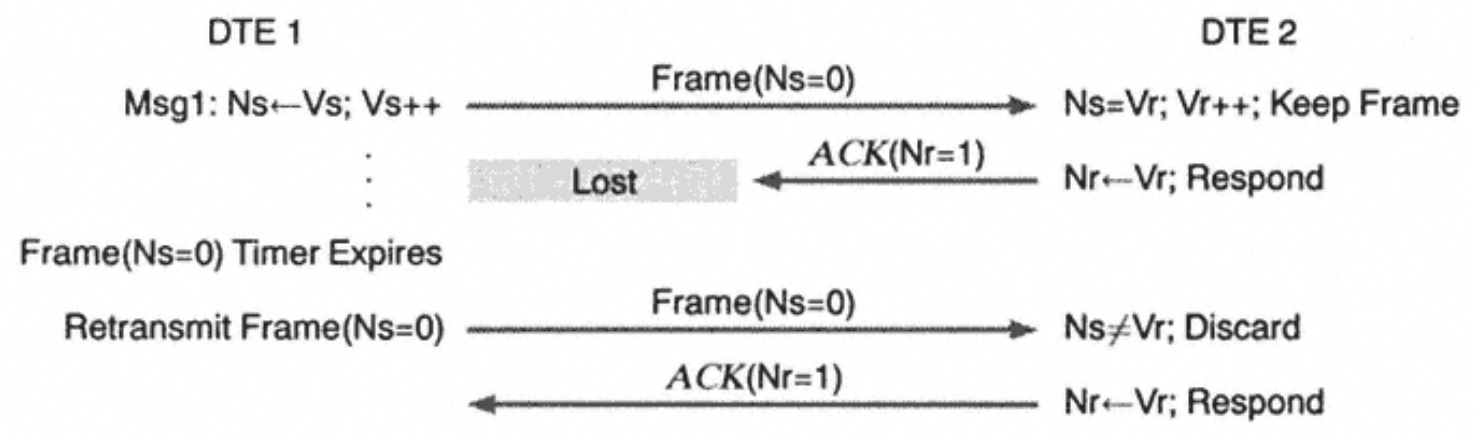

Figure 6.35
Acknowledging (and discarding) a duplicate frame.

First, the transmitting DTE should be operated in such a way as to keep the channel "full".

Second, acknowledgments must be returned at regular intervals in order to minimize the number of retransmissions required if one of the previous frames is received in error.

In the stop-and-wait protocol, two sequence numbers are neededone for the previous frame and one for the current frame. The sequence numbers required in a go-back-N protocol is not so readily apparent. If only two sequence numbers are used (i.e., 0 and 1, as in stop-and-wait), the number of frames that the transmitting DTE can send without receiving an acknowledgment is one, as Figure 6.36 illustrates (note that the increment operator, ++, is assumed to operate modulo-2, meaning that the possible values for Vs and Vr in this example are 0 and 1).

In Figure 6.36, the receipt of *ACK*(Nr = 0) (by DTE 1) can mean *either* of the following:

1. Two frames were received (Ns = 0 and Ns = 1), and DTE 2 is expecting a third (to be sent with sequence number Ns = 0).

2. DTE 2 is still expecting the first frame (that is, frame Ns = 0 was lost), meaning that both frames must be retransmitted.

Since it is necessary for protocols to be well defined and unambiguous, the number of unacknowledged frames must be less than the total number of possible sequence numbers (otherwise, problems such as those described

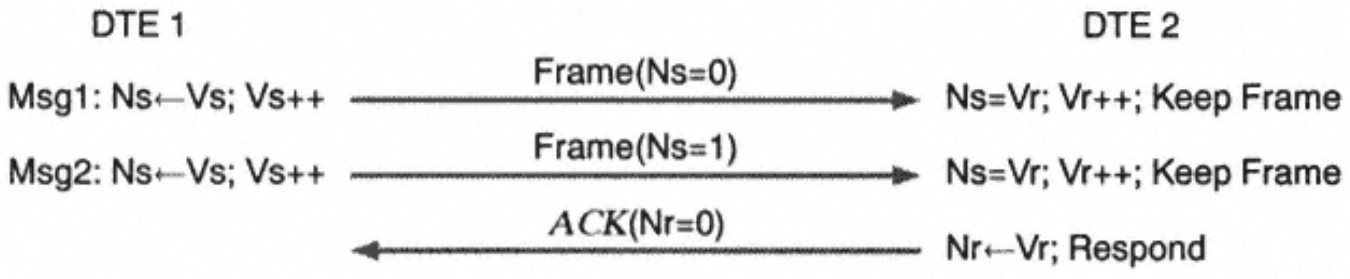

Figure 6.36
A problem with sequence numbers.

above can occur). In short, if there are *N* sequence numbers available, then at most *N* 1 frames can be unacknowledged.

Consider, for example, a protocol that supports three sequence numbers (0, 1, and 2); the transmitting DTE must wait for an acknowledgment after sending two frames (each with a unique sequence number). Upon receipt of the acknowledgment, further frames may be transmitted (or retransmitted should a negative acknowledgment be received). Figure 6.37 illustrates a typical frameacknowledgment cycle ("Next" denotes the next pair of send sequence numbers that DTE 2 is expecting).

Protocols that allow a transmitting DTE to send a number of frames before the receiving DTE responds are referred to as sliding window protocols. The term "sliding window" can have a number of connotations:

As each message is sent, the transmission window (i.e., the number of frames that can be sent before an acknowledgment is required) is decreased in size by 1; transmissions continue until the window has a size of 0, at which point they cease. Whenever an acknowledgment is received, the transmission window is opened (the exact amount depends upon the value of Nr returned by the receiver); transmissions can resume, decreasing the window size.

In the above example, DTE 1 can send two frames before waiting for a reply; it therefore has a transmission window of size 2. Each acknowledgment received indicates that the two messages sent were received correctly, fully opening the transmission window.

The window "slides" forward, indicating the sequence numbers to be used for the next set of message(s) to be sent. In Figure 6.37, three sequence numbers exist (0, 1, and 2); since they are always being re-used, the list of

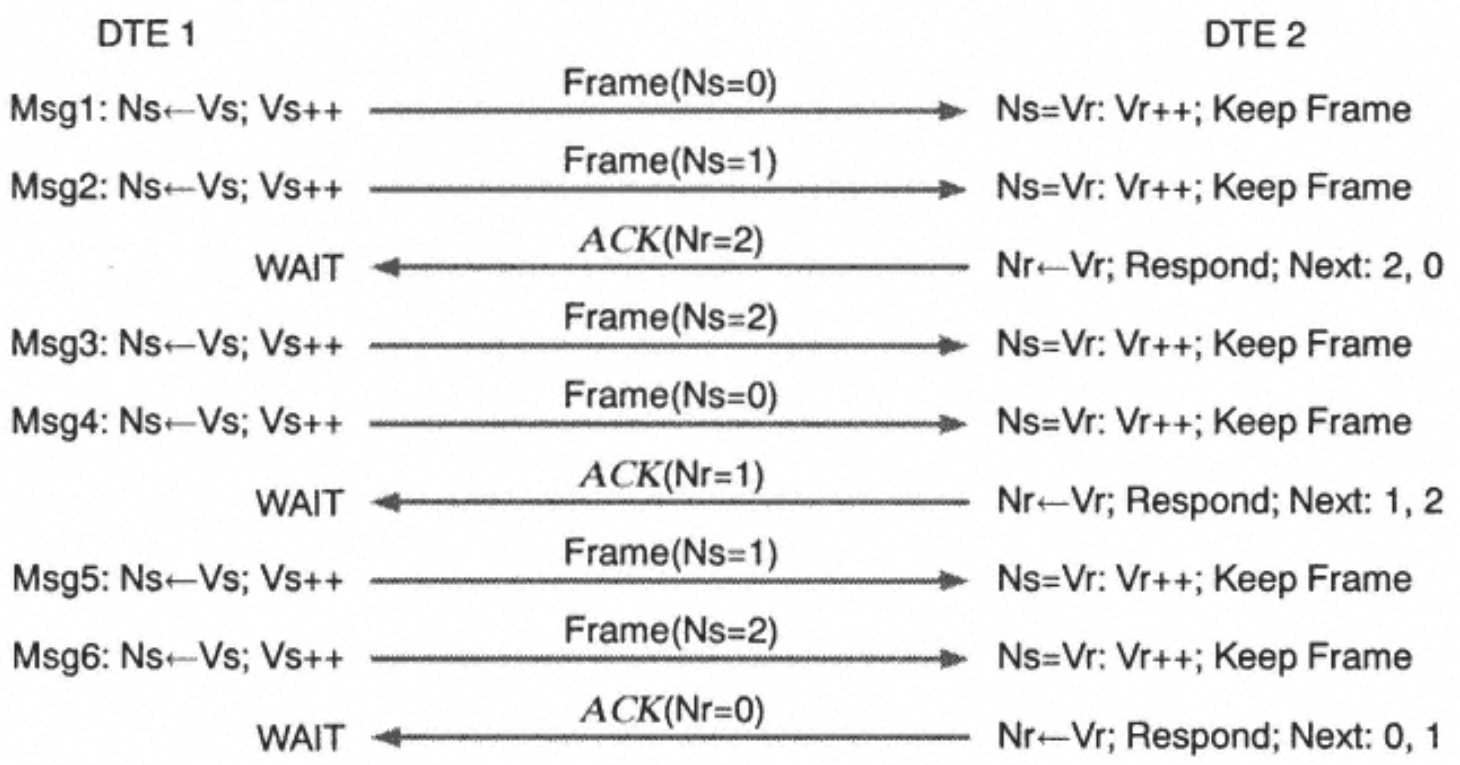

Figure 6.37
Go-back-N using three sequence numbers.

available sequence numbers can be considered as a long repeating sequence of the same three values:

0 1 2 0 1 2 0 1 2 ...

At any moment, two of the sequence numbers can be in use; as acknowledgments are received, the window slides over the next set of possible sequence numbers. For example, after the initial transmission, the window sits over 0 and 1:

[0 1] 2 0 1 2 0 1 2 ...

After the first acknowledgment, the window slides over the next two sequence numbers:

0 1 [2 0] 1 2 0 1 2 ...

The number of frames that can be sent without being acknowledged is known as the window size. Initially, the transmission window is set to the value of the window size, but as each frame is sent, it is decreased until it reaches 0 (i.e., the window is closed), at which point no further frames can be sent. In the example described above, the window size is 2.

To maximize throughput, the transmission window should never be allowed to close, which means that the receiving DTE must respond with sufficient acknowledgments to ensure an uninterrupted flow of frames. This can be achieved in a number of ways: every frame can be acknowledged, or the window size can be made large enough to require fewer acknowledgments. In Figure 6.38, there are five possible sequence numbers (giving a maximum transmission window size of 4). By responding after every third frame, the

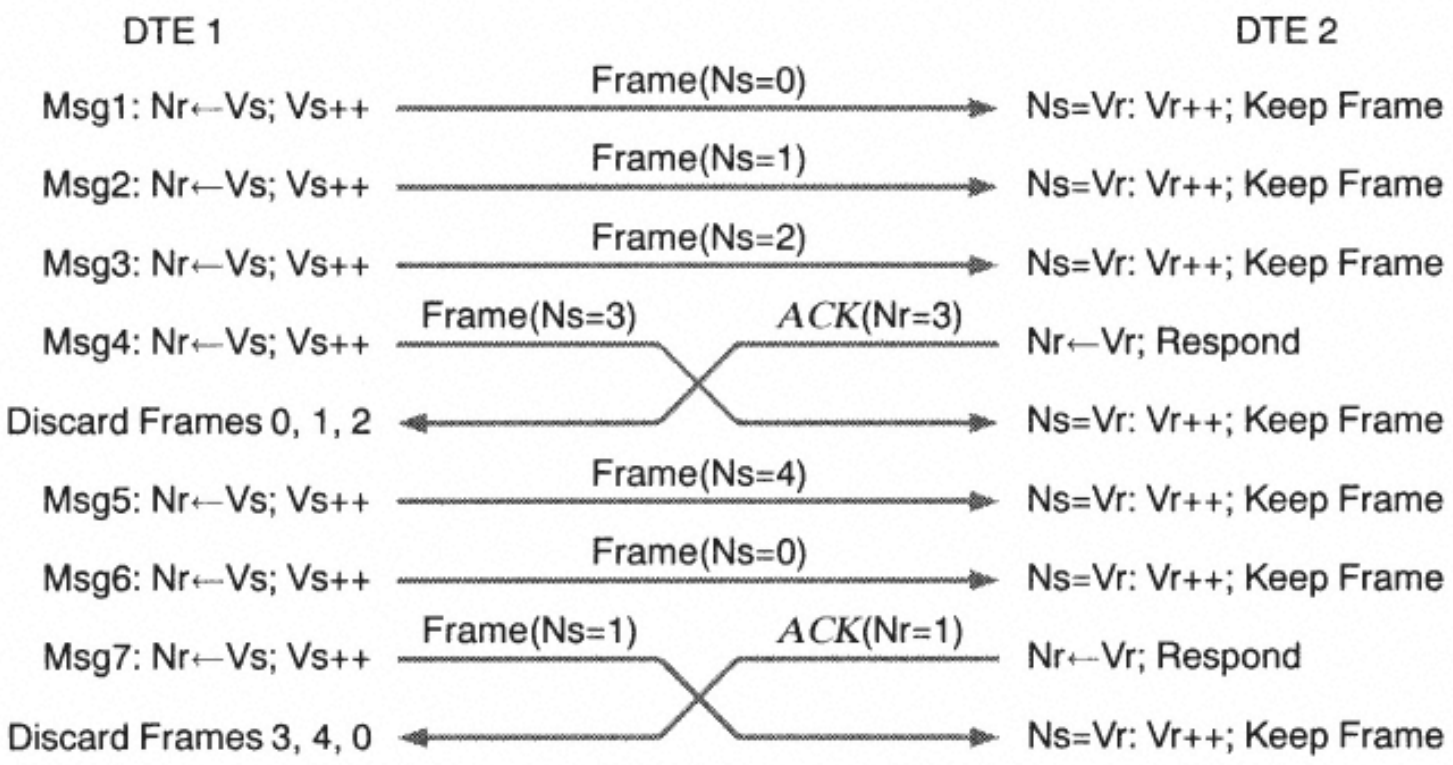

Figure 6.38
With a window size of 4, DTE 2 can respond every three frames.

receiving DTE can keep the window open (if it had waited until the fourth frame, the transmission window would have closed).

As with so many things in data communications, there are tradeoffs that must be considered when specifying a window size. For example, the larger the window, the more pending-frame space required (since a frame cannot be discarded until it has been acknowledged); furthermore, if a frame is *NAK*ed, all subsequent pending frames must be retransmitted. With these things in mind, many systems use a window size of 2, which is similar to the double buffering used in the UART: one frame is being transmitted while the other is being prepared.

The number of frames sent by the transmitting DTE can be controlled by the receiving DTE imposing flow control. Flow control restricts the rate at which frames are acknowledged (in the most extreme case, if the receiving DTE does not respond at all, the transmitting DTE is forced to retransmit all pending frames, as their timers expire). As shown in Figure 6.39, given a window size of 2, the receiving DTE could turn the go-back-N protocol into a stop-and-wait protocol (assume three sequence numbers).

In Figure 6.39, the DTE 2 never allows DTE 1 to send more than one frame before the transmission window closes. This effect is achieved by having DTE 2 respond with an acknowledgment Nr value equal to the value of Ns taken from the most recently received frame (as opposed to the value of Vr, the expected sequence number of the next frame). For example, after receiving frames Ns = 0 and Ns = 1, instead of responding with Nr = 2 (an indication that frames Ns = 0 and Ns = 1 were received correctly), the receiving DTE responds with Nr = 1, thereby allowing the window to move forward one sequence number only, as shown in Figure 6.40. To avoid DTE 1 retransmit-

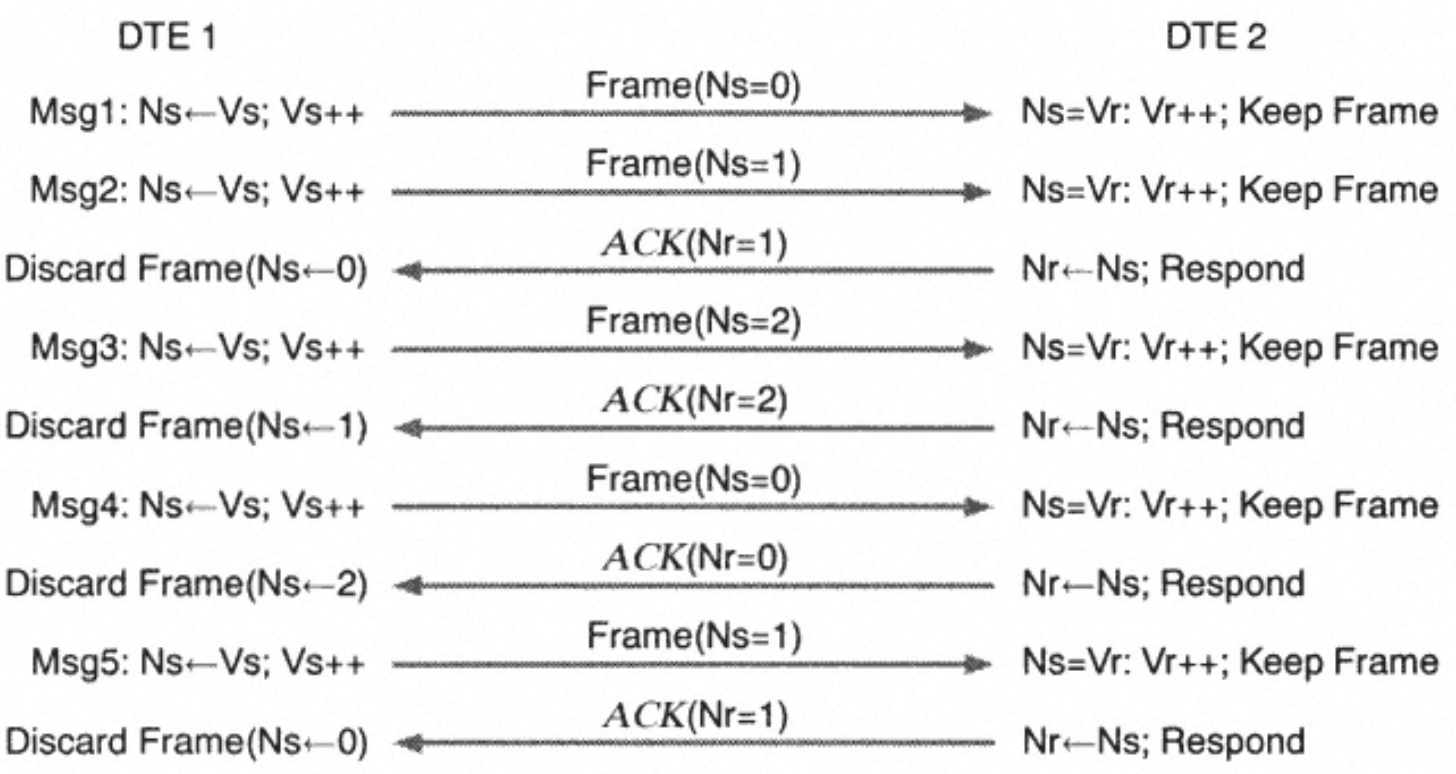

Figure 6.39
A go-back-N implementation of stop-and-wait.

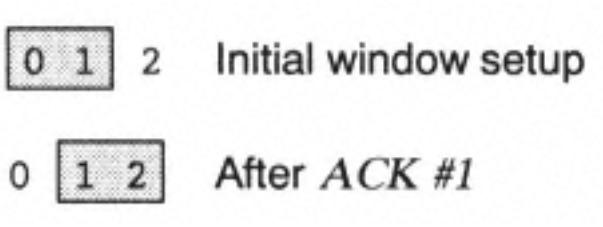

Figure 6.40
The transmission window for Figure 6.39.

ting any pending frames, DTE 2's acknowledgments must arrive before each pending frame's timer expires.

The sequence numbers available in any protocol usually equal a power of 2 (if the sequence numbers are sent as a binary number) and depend upon the protocol. Most protocols vary the number of possible sequence numbers from 2 (a 1-bit sequence number, 21) to 8 (a 3-bit sequence number, 23). In situations where there may be long transmission delays, up to 128 sequence numbers can be specified (a 7-bit sequence number, 27).

Until now, data transfer has been half-duplex; that is, one DTE has been transmitting frames and the other has been receiving them. Should both DTEs have information to exchange, the data transfer methods examined thus far will force one station to wait until the other has completed its transfer. Since the communication channel is full-duplex, it would appear that at the channel level, at least, there is no reason that information cannot be sent in both directions simultaneously using a bidirectional data transfer.

To accomplish this transfer, the change required to the go-back-N protocol is actually quite simple: acknowledgments must be sent with the frames using a technique known as piggybacking. In piggybacking, each frame is sent with its own send sequence number (i.e., Ns) *and* the sequence number of the frame next expected to be sent by the destination DTE (i.e., Nr), obtained from the DTE's receive state variable, Vr). The frame still carries a message.

When a DTE receives a frame, it must check the values of both Ns and Nr:

Ns. As before, the value of Ns is compared to the DTE's receive state variable Vr. If the two are equal, the frame is accepted and Vr is incremented. However, if the values are not equal, the frame is discarded.

Nr. The value of Nr is an acknowledgment, informing the local DTE of the frames that have been accepted by the remote DTE. A value of *N* acknowledges all frames up to frame *N* 1; the acknowledged frames can be discarded by the local DTE.

A possible exchange between two DTEs performing a bidirectional transfer is presented in Figure 6.41. In this example, DTE 2, rather than responding with an acknowledgment, responds with a frame containing the value of Nr = 1, indicating that all frames with sequence numbers less than 1 have been acknowledged. Since there is only one frame format, the original frame sent from DTE 1 has a value of Nr = 0.

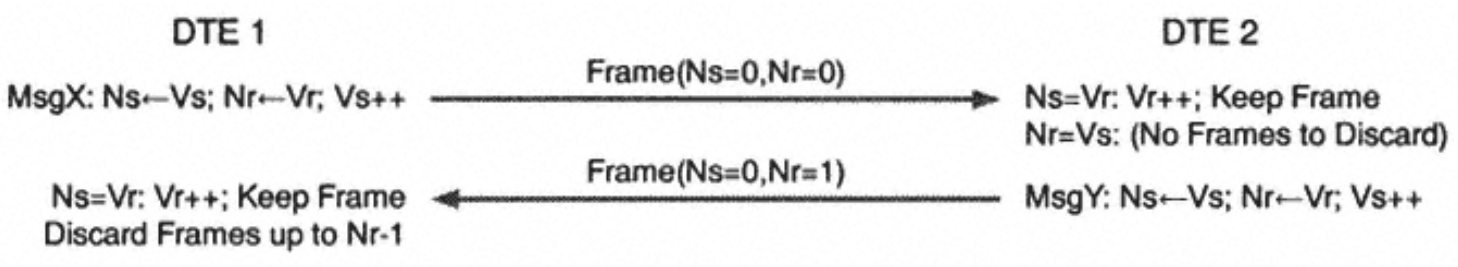

Figure 6.41
Bidirectional transfer with piggybacked acknowledgments.

If a station does not have information to send but has a frame to acknowledge, it responds with an acknowledgment rather than a frame. The acknowledgment contains the value of the station's next expected sequence number (i.e., Nr). For example, in Figure 6.42 DTE 2 has a single frame to transmit, which can be used to carry the value of Nr; all subsequent responses must be sent as acknowledgments. Whenever a frame or an acknowledgment is received, the value of Nr is examined: all pending frames with Ns values less than that of Nr are discarded. Furthermore, the value of Nr sent in the frame remains constant until the value of Vr changes. Note that in the go-back-N protocol, only frames are acknowledged, acknowledgments are not.

If a frame is lost or received in error, the same rules apply in a bidirectional transfer as in a unidirectional half-duplex: the receiving DTE must inform the transmitting DTE of the error with a *NAK*; or, if frames have not been received, the transmitting DTE should time out and retransmit the frames not yet acknowledged. For example, a frame received out of sequence (because of the loss of a previous one), should result in a negative acknowledgment, as illustrated in Figure 6.43. Note that when the frame is being retransmitted, the value of Ns cannot change (since this is the frame's send sequence number); however, the value of Nr can change if any frames are received and must be acknowledged.

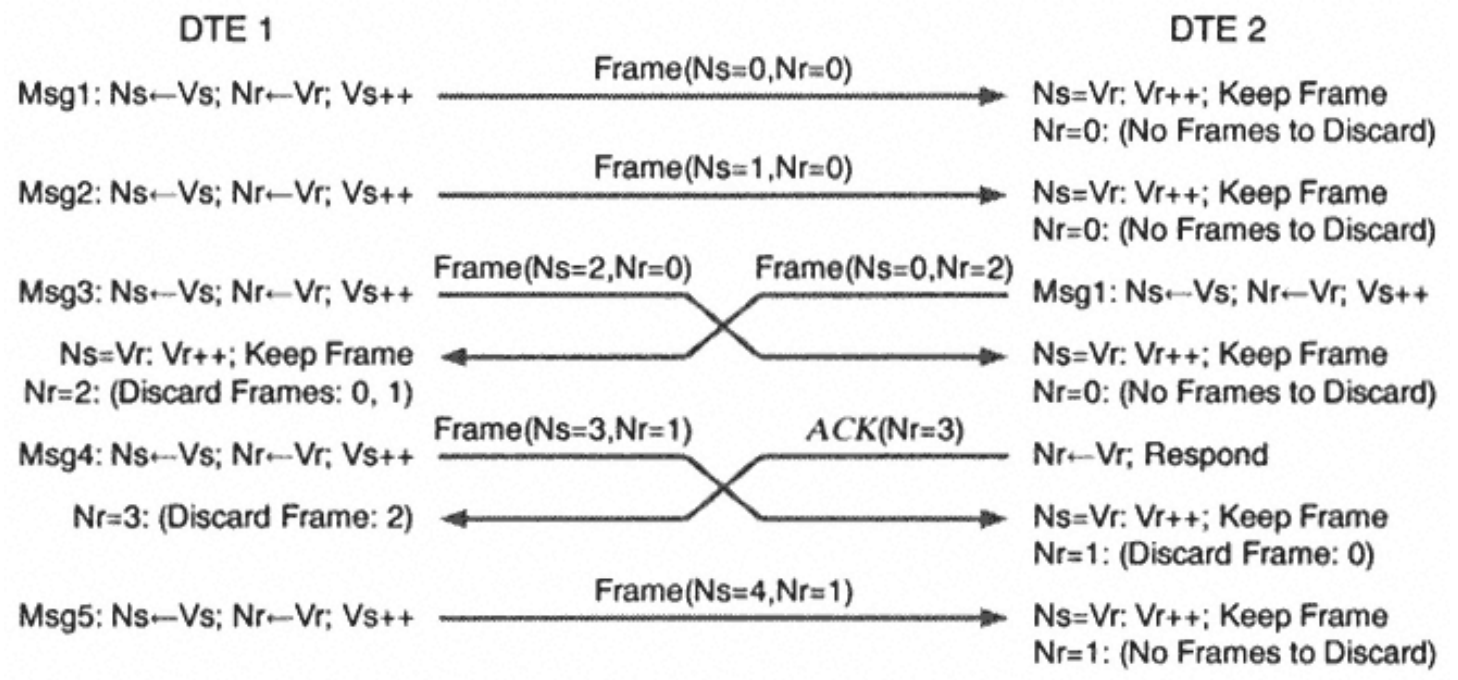

Figure 6.42
Mixing frames and acknowledgments.

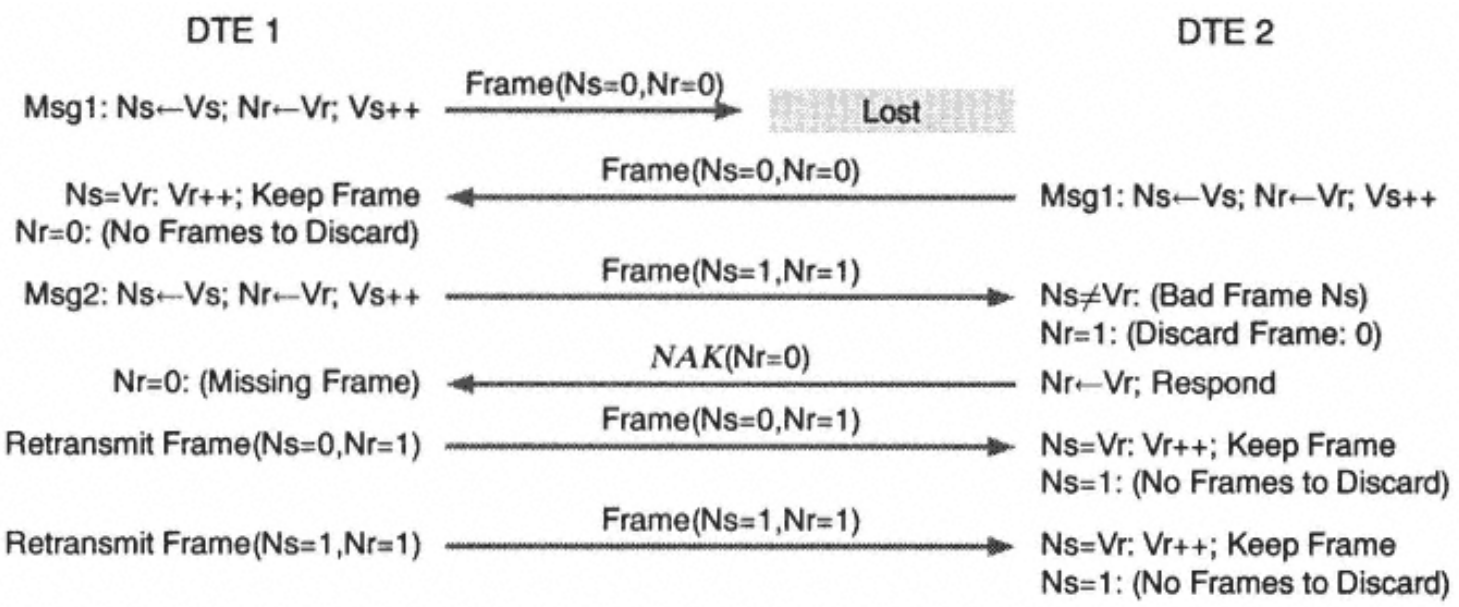

Figure 6.43
Error handling in a bidirectional transfer.

Similarly, unacknowledged frames should also be retransmitted once the timer associated with each frame has expired; in Figure 6.44, two frames are lost but are retransmitted once the timer has expired.

Selective Retransmission

The go-back-N protocol offers considerably more flexibility and throughput than the stop-and-wait protocol. However, go-back-N has a problem in that when a frame is received out of sequence or in error, all frames subsequently received are discarded until the frame with the correct sequence number is retransmitted. Consider the situation shown in Figure 6.45 in which frames Ns = 2 and Ns = 3 are retransmitted, even though only frame Ns = 1 was in error. Ideally, only frames received in error should be retransmitted, to avoid unnecessary retransmissions.

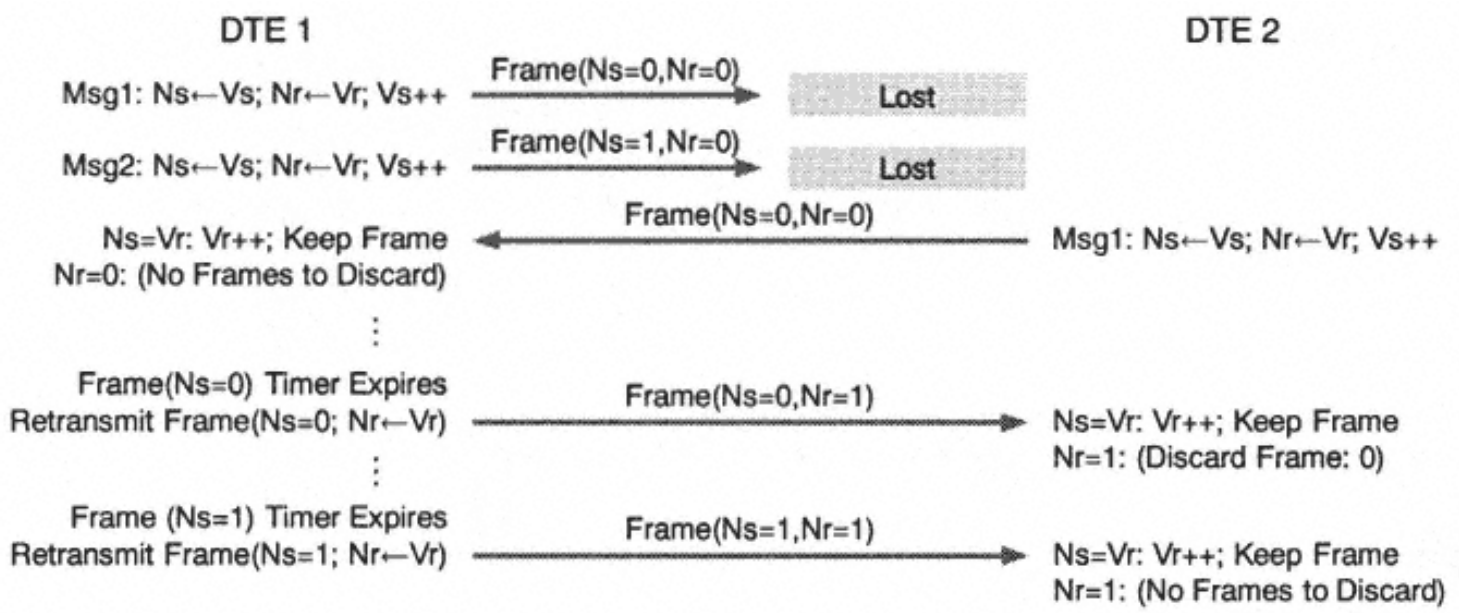

Figure 6.44
Frame retransmission after timeout.

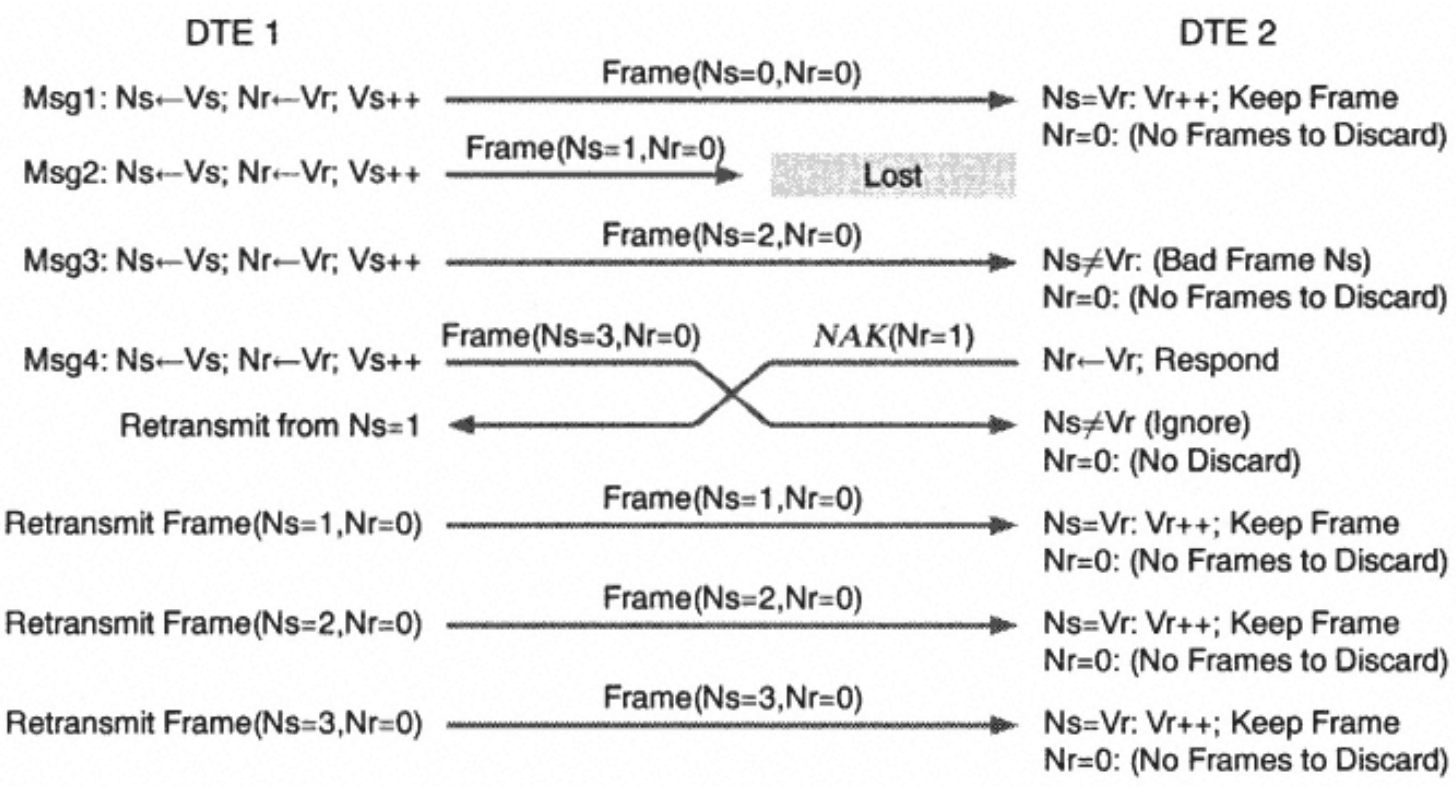

Figure 6.45
A go-back-N limitation: the retransmission of valid frames.

Selective retransmission is a technique in which only frames that have not been received correctly (or have apparently not been received) are retransmitted. In selective retransmission, the receiving DTE explicitly acknowledges each frame as it is received. A missing frame will not be acknowledged, causing the transmitting DTE to eventually retransmit it (either by detecting it or by the time limit associated with the frame expiring). See Figure 6.46 for an example.

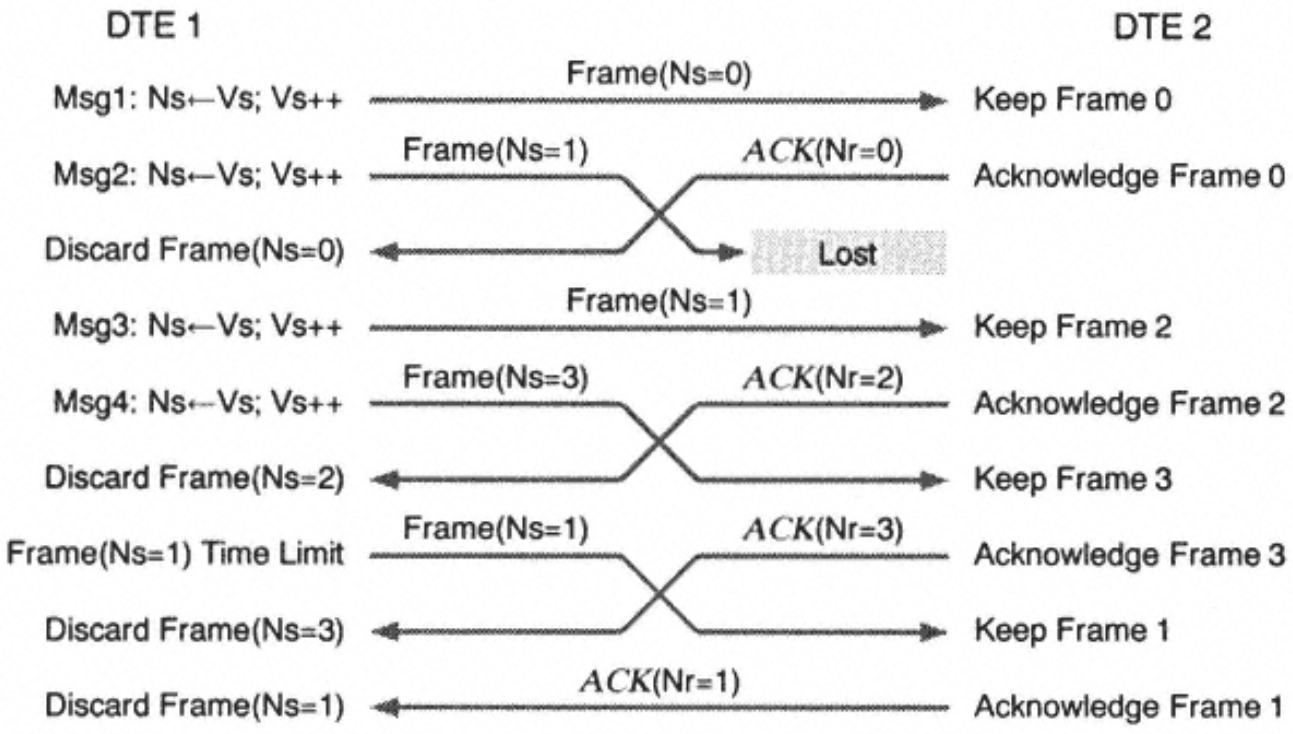

Figure 6.46
Selective retransmission: only one frame is retransmitted.

The time it takes to recover from the loss of a frame depends upon the time limit associated with each frame. This delay can be reduced by having the receiving DTE transmit a *NAK* that explicitly identifies the missing frame. In Figure 6.46, DTE 2 could transmit a *NAK* 1 as soon as frame Ns = 1 is determined to be missing (i.e., after frame Ns = 2 is received). Upon receipt, of the *NAK*. DTE 1 can retransmit frame Ns = 1.

The loss of an acknowledgment can lead to some interesting problems if there are limited sequence numbers. Consider the situation shown in Figure 6.47, in which two sequence numbers are used to represent the frames sent.

When DTE 2 receives the retransmitted frames Ns = 0 and Ns = 1, there is no indication that these are retransmissions. The result is that DTE 2 will keep *two* copies of frames Ns = 0 and Ns = l. The concept of a transmission window that slides as each acknowledgment is received will not work for selective retransmission. For example, the process of using three sequence numbers (0, 1, and 2), requiring the transmitting DTE to have at most two frames waiting for acknowledgment, and sliding the window forward as acknowledgments are received, fails if frames Ns = 0 and Ns = 1 are transmitted and an acknowledgment is received for frame Ns = 1 only. If the transmitting DTE sends frame Ns = 2 and retransmits frame Ns = 0, the receiving DTE once again has no indication that frame Ns = 0 is a retransmission.

A number of solutions to this problem exist. In one, a form of sliding window is used that moves forward only when the "oldest" frame has been acknowledged.

Although selective retransmission may appear to offer numerous advantages, it is not widely used. The main problem with it is the storage overheads required by the receiving DTE for holding frames while waiting for any intermediate missing frames.

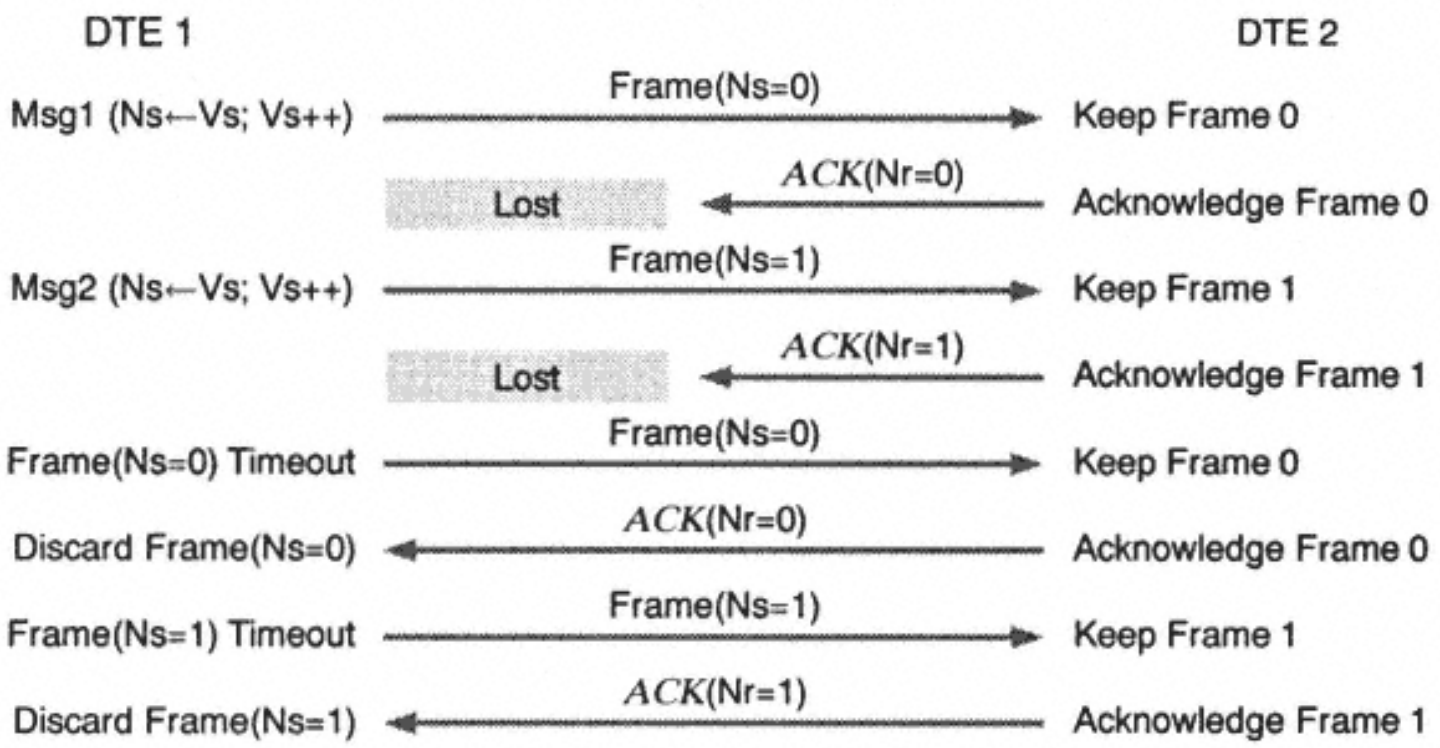

Figure 6.47
Sequence numbering problem with selective retransmission.

6.4.2
Forward Error Recovery

In certain situations, it may not be possible or even desirable to retransmit a frame that has been received in errorspace probes with limited storage capacity that are hundreds of millions of miles from Earth cannot be expected to hold onto frames while waiting for acknowledgments. Solutions such as transmitting the same message twice not only take twice the bandwidth but, as shown in Section 6.3, do not offer the receiving DTE any indication as to which frame (if any) is correct.

In applications such as the one just described, it is necessary to supply the receiving DTE with the capability to both detect and, if possible, *correct* any errors found in the frame. This is known as forward error recovery and requires the transmitting DTE to send sufficient additional information within the frame to allow the receiving DTE to correct the detected errors. The events that occur in forward error recovery are illustrated in Figure 6.48.

The Hamming code is one of several error-correcting codes that use parity bits to both detect and correct errors. In the Hamming code, a frame consists of two types of bit: data (obtained from the data to be transmitted), and parity (generated from the data bits stored in the frame). The parity bits are stored in those frame-bit positions that are powers of 2; for example, 1, 2, 4, 8, and so on (the rightmost bit of a frame is considered to be bit 1); the data bits are stored in the remaining frame bits.

The value of a parity bit is determined by its position in the frame; the parity bit stored in bit 2*i* includes all frame-bit positions that have bit 2*i* set. Thus, the parity bit stored in frame-bit position 20 (i.e., frame bit 1) is generated from frame bits 3, 5, 7, 9, and so on (since these frame-bit positions have bit 1 set). Table 6.1 shows the association of frame bits and parity bits in a 7-bit frame.

The 7-bit frame shown in Table 6.1 has three parity bits (stored in framebit positions 1, 2, and 4, since these are powers of 2) and allows four data bits (stored in the remaining frame-bit positions, 3, 5, 6, and 7). The Hamming code uses even parity, meaning that the values of the parity bits can be obtained using exclusive-or; Table 6.2 shows how the three parity bit values in Table 6.1 are generated (Px denotes the parity bit in frame-bit position x, while Dx denotes the value of the data bit in frame-bit position x).

For example, to transmit the nibble `1101` using the above algorithm first

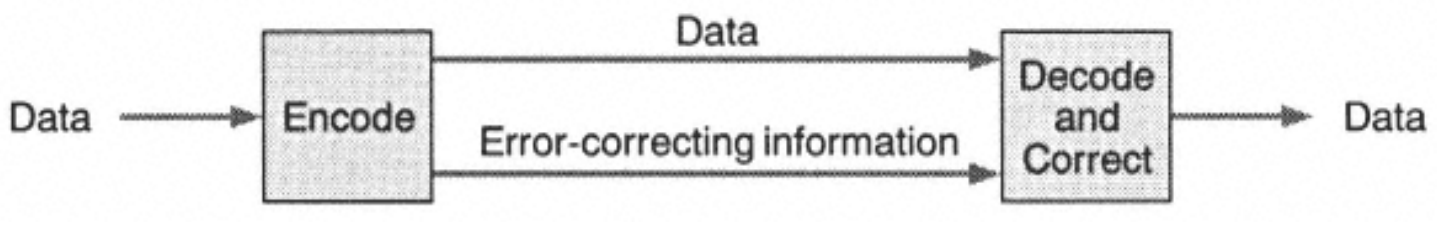

Figure 6.48
Forward error recovery.

Table 6.1 The association of frame bits and parity bits.

Frame bit	Binary value	Parity bits		
		4	2	1
1	001			X
2	010		X	
3	011		X	X
4	100	X		
5	101	X		X
6	110	X	X	
7	111	X	X	X

Table 6.2 Parity-bit values for Table 6.1.

P1 = D3 ⊕ D5 ⊕ D7

P2 = D3 ⊕ D6 ⊕ D7

P4 = D5 ⊕ D6 ⊕ D7

requires the assignment of the data bits; in this case D3, D6, and D7 are set while D5 is clear. The assignment of the bits in a 7-bit frame is shown in Figure 6.49.

To obtain the parity bits, it is necessary to generate each parity from the data bits in each frame bit. The parity bits are therefore assigned as follows:

P1 = 1 ⊕ 0 ⊕ 1
P2 = 1 ⊕ 1 ⊕ 1
P4 = 0 ⊕ 1 ⊕ 1

(i.e., P1 is clear, P2 is set, and P4 is clear). The frame to be transmitted is assembled from the four data bits and the three parity bits, and is shown in Figure 6.50.

Once the data and parity bits have been assigned to the frame, the frame can be transmitted, and upon receipt of the frame the parity and data bits can be extracted; however, the question remains: is the frame in error, and, if so, how can the error be corrected?

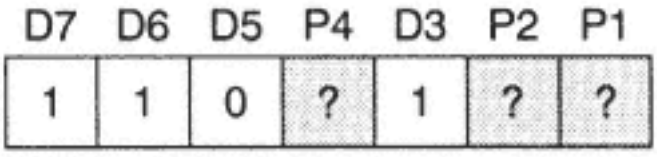

Figure 6.49
Frame including data and parity bits.

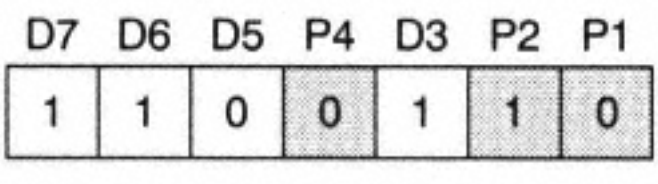

Figure 6.50
Frame including data and parity bits.

In Chapter 3, error detection with parity bits was achieved by calculating the parity on the entire frame (including both the data bits and the parity bit). The same approach is taken with the Hamming code, but here, since there are several parity bits, each parity must be calculated. In the 7-bit frame example, there are three parity bits, which means that three parity checks are required, as shown in Table 6.3 (note that Ex denotes the status of the parity associated with frame-bit *x*; 0 means that no error was detected, while 1 signals an error).

In this example, the receipt of the frame `1100110` results in

E1 = 0 ⊕ 1 ⊕ 0 ⊕ 1
E2 = 1 ⊕ 1 ⊕ 1 ⊕ 1
E4 = 0 ⊕ 0 ⊕ 1 ⊕ 1

where E1, E2, and E4 are all clear, indicating that no error was detected, meaning that the nibble value is `1101`.

Should an error occur (remember, this code only detects a single-bit error), what happens? Assuming that the same frame was transmitted but that instead of `1100110` the frame `1110110` was received (i.e., bit D5 was inverted), the result of the error check will be

E1 = 0 ⊕ 1 ⊕ 1 ⊕ 1
E2 = 1 ⊕ 1 ⊕ 1 ⊕ 1
E4 = 0 ⊕ 1 ⊕ 1 ⊕ 1

In this situation, E1 and E4 are set while E2 is clear. Obviously, the error has been detected, but how can it be corrected?

To understand how error correction is achieved it is necessary to reexamine Table 6.1. Each frame bit in this table is associated with a unique encoding of parity bits; for example, frame bit 6 is represented by parity bits P4 and P2 being set, while P1 is clear. Should any frame bit be inverted in the transmission, the parity check at the receiver will be incorrect; furthermore,

Table 6.3 Hamming code parity checking.

E1 = P1 ⊕ D3 ⊕ D5 ⊕ D7

E2 = P2 ⊕ D3 ⊕ D6 ⊕ D7

E4 = P4 ⊕ D5 ⊕ D6 ⊕ D7

only parity values associated with the inverted frame bit will signal an error. For example, if frame bit 6 was inverted, parity bits P4 and P2 would be set while P1 would be clearindicating that frame bit 6 (from 4 × 1 + 2 × 1 + 0) was in error and should be inverted.

Therefore, when the frame `1110110` was received, the result of the error check indicated E1 and E4 set and E2 clear. This means that 4 × 1 + 1 × 1, or frame bit 5, was inverted (i.e., D5 should have a value of 0 rather than 1). By inverting bit 5 in the frame (i.e., `1100110`), the corrected nibble can be extracted as `1101`.

The Hamming code examined in this section will detect any single-bit errors, including errors in the parity bits. Single- and multiple-errorcorrecting Hamming codes are often used when storing data in both internal and external storage. Other examples of error correcting with the Hamming code are given as exercises.

6.5
Synchronous Protocols

Broadly speaking, synchronous protocols are divided into two categories: byte-oriented and bit-oriented.

6.5.1
Byte-Oriented Synchronous Protocols

Byte-oriented synchronous protocols are those that use the control characters found in character codes (such as ASCII or EBCDIC) to delimit frames and support other aspects of the protocol. Although numerous byte-oriented protocols exist, the following is a brief introduction to the "grandparent" of them all: IBM's early (and highly successful) Binary Synchronous Communications Protocol (known as Bisync or BSCso named because it allows the synchronous transmission of binary data). Bisync was designed around the EBCDIC character set, although ASCII implementations do exist.

The Bisync protocol is intended for batch processing, in which users submit their jobs to a remote job entry (or RJE) station, typically an input device and a line printer. The job, consisting of a file of instructions, is read and transmitted to the central host for processing. Once the processing has finished, the output is spooled to disk and, when the opportunity arises, sent to the user's RJE for printing.

Bisync is a stop-and-wait protocol that allows data transfer between pairs of stations. It has a window size of 1 (i.e., there can be at most one pending data frame) and two sequence numbers (0 and 1). Each data frame is associated with a sequence number: the frames transmitted first, third, fifth, and so on, with sequence number 1; and the frames transmitted second, fourth, sixth, and so on, with sequence number 0. Data frames do not carry the sequence numbers; instead, sequence numbers are returned with acknowledgment frames. Acknowledgment frames are transmitted as two characters,

Table 6.4 Bisync acknowledgment sequences.

Acknowledgment	EBCDIC	ASCII
ACK 0	*DLE* `0x70`	*DLE* 0
ACK 1	*DLE/*	*DLE* 1

as shown in Table 6.4 (note that the sequences are different for ASCII and EBCDIC).

Whenever a data frame is received correctly, the receiving station changes the acknowledgment (from 0 to 1 or vice versa) and responds with it. Once a data frame has been properly acknowledged, the transmitting station changes the expected acknowledgment (from 0 to 1 or vice versa). A typical transmission is shown in Figure 6.51.

Any data frame received in error is negatively acknowledged (using a *NAK*). The transmitting station is expected to retransmit the frame, while the receiving station is to reply with the correct acknowledgment. An example of frame retransmission is given in Figure 6.52 (in this example, assume that the acknowledgment for frame *N* is *ACK* 1).

If a frame is not acknowledged (i.e., the receiving station never receives the frame or the acknowledgment is lost) within 2 seconds, the transmitting station must poll the receiving station to determine the value of the last acknowledgment sent. Polling takes place using the *ENQ* (enquire) control character. There are three possible outcomes of the poll:

The correct acknowledgment is received. The frame was received correctly by the receiving station, but the acknowledgment was lost. As shown in Figure 6.53, the next frame can be transmitted at this point (assume that frame *N* is to be acknowledged with *ACK* 0).

The wrong acknowledgment is received. The receiving station never received the data frame and is still responding with the acknowledgment of the

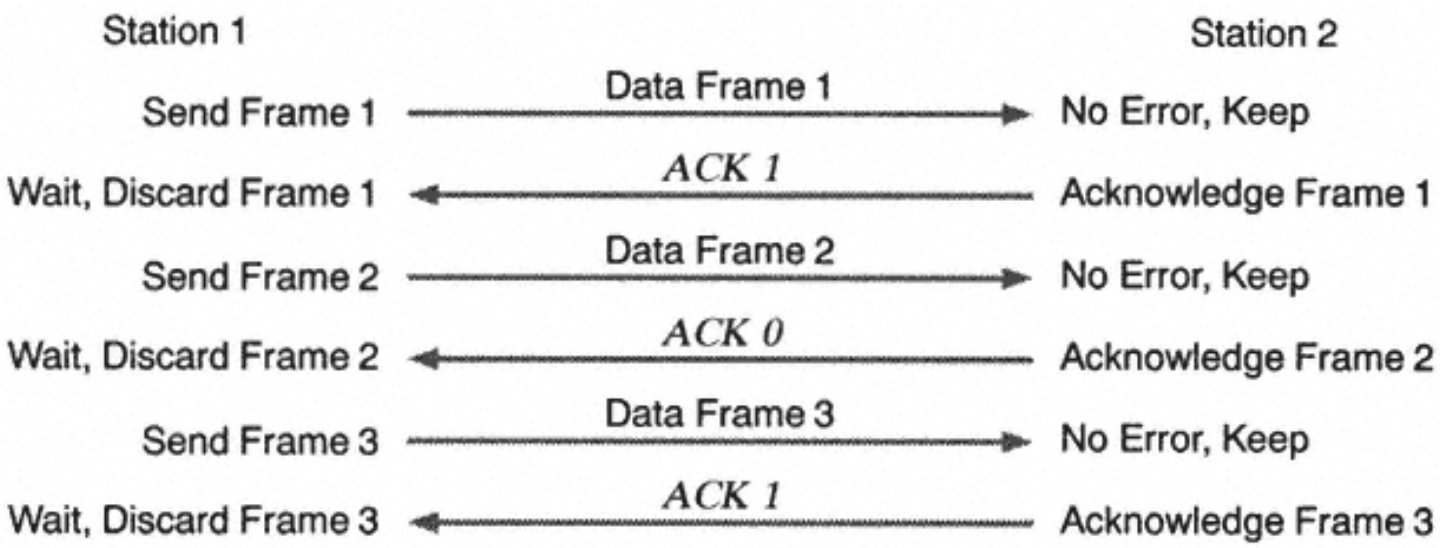

Figure 6.51
A typical error-free Bisync sequence.

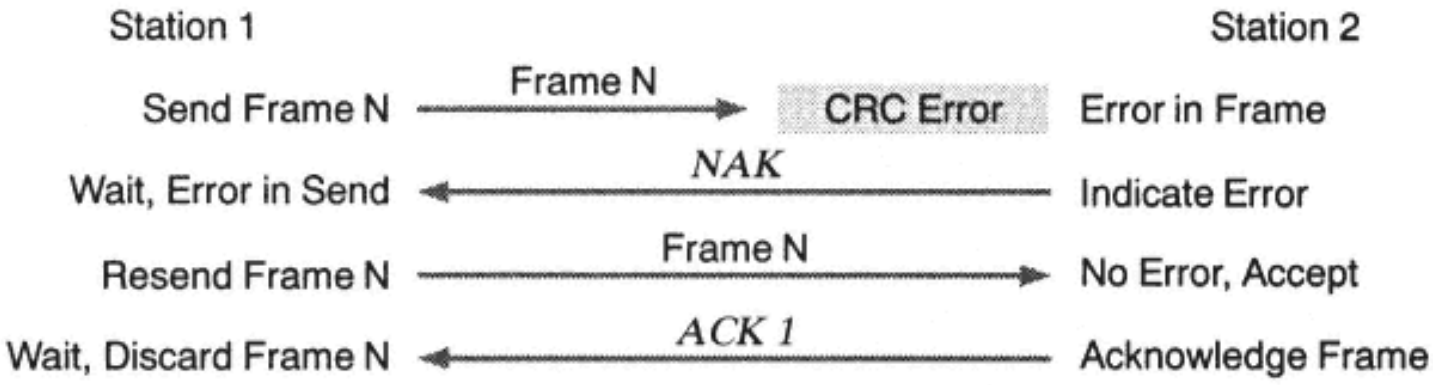

Figure 6.52
Frame transmission due to CRC error detection.

last correctly received frame. The transmitting station must retransmit the last data frame, as shown in Figure 6.54.

The receiving station never responds. It is possible that the receiving station is momentarily off-line or simply too busy to respond to the poll. To accommodate situations such as these, the transmitting station will wait for 3 seconds and poll again. If, after three such cycles, no response is obtained, the transmitting station aborts the transmission.

There are a number of variations and exceptions to the time and number of repeats the transmitting station performs. In some installations, for example, the number of polls can be greater to accommodate transmission delays, or the time between polls can be increased.

Other responses. There are other responses allowed by Bisync. For example, should the receiving station be unable to accept more data frames because of problems such as no paper in the line printer, the communication channel can be kept active by the receiving station responding with *WACK*s (wait acknowledgments). The transmitting station continues polling (with *ENQ*s) until an acknowledgment is received.

Although there are exceptions (for example, a station with only a line printer or a station with only a card reader), most Bisync implementations allow either station to transmit and receive information. Bisync requires the station about to transmit to send an *ENQ* character, and the potential receiving station is to respond with an *ACK* 0, after which the data transfer may proceed. In Figure 6.55, station 1 is bidding for the line.

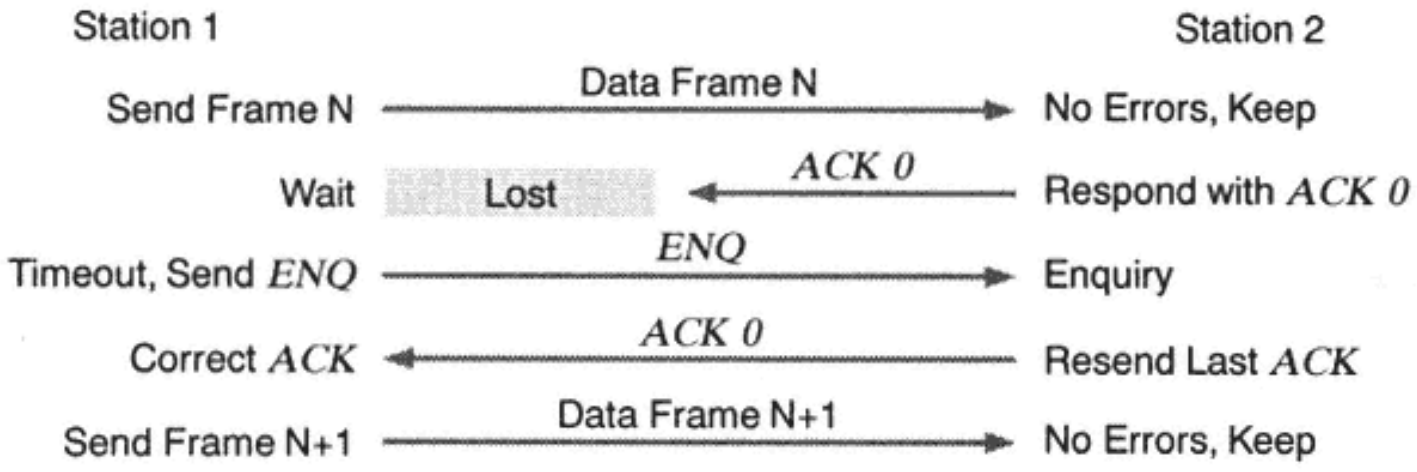

Figure 6.53
The correct acknowledgment signals transmission of the next frame.

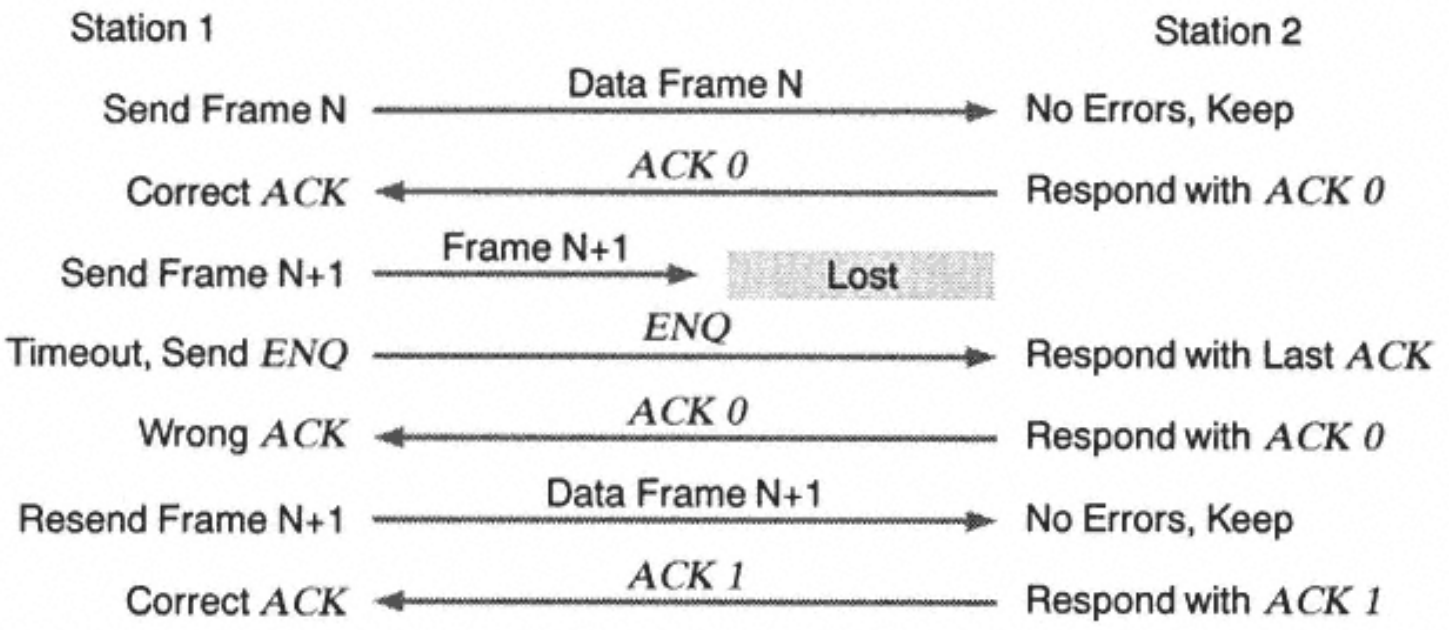

Figure 6.54
The previous acknowledgment signals a retransmission.

One problem to consider with Bisync implementations is what happens if both stations attempt to transmit at the same time. From the algorithm described above, both will send *ENQ*s and wait for the response. Since only one station can transmit at any time (remember, Bisync is half-duplex), one of the two stations must "win" the bid. The solution adopted by Bisync is to designate one station the master and the other the slave. A slave station must defer transmission to the master station whenever a conflict occurs. When both stations bid for the channel (with *ENQ*s), the master station is guaranteed to win the bid because it has a shorter polling wait period: one second as opposed to three seconds for the slave.

Once a station has finished transmitting, the receiving station is signaled with an end-of-transmission character, or *EOT*. At this point, both stations return to the state where either can bid for the channel. *EOT* can also be used to abort a transmission. If the transmitting station sends an *EOT*. the receiving station returns to the line-bidding state.

Although Bisync is half-duplex, there is a provision for forcing the transmitting station to give up the line, thereby allowing the receiving station to send a message. Either station, when receiving, can respond with an *RVI*. or

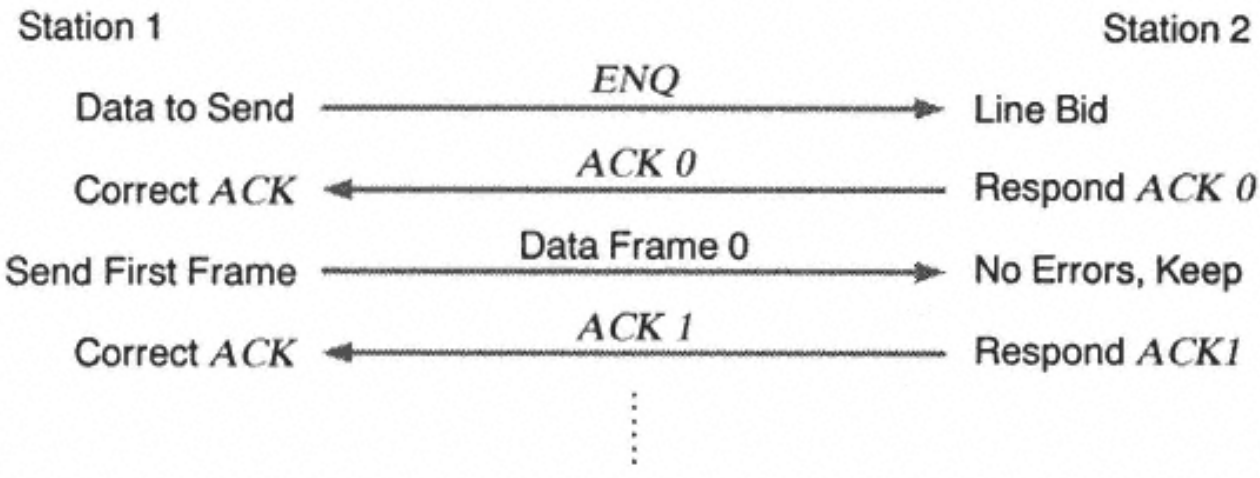

Figure 6.55
The line-bid sequence.

reverse interrupt. which is equivalent to an acknowledgment. Upon receipt of the *RVI*. the transmitting station sends one more block, waits for the acknowledgment, and sends an *EOT*. The receiving station can then bid for the line.

The Bisync Frame

Bisync, like any synchronous protocol, requires the use of framing information to allow a receiving DTE to distinguish between data and the idle channel. It uses special control charactersfor example, start-of-frame is indicated by either of the following:

*SOH*Start of Header.

*STX*Start of TeXt.

End-of-frame is also a control character, represented by one of

*ITB*Intermediate Transmission Block.

*ETB*End of Transmission Block.

*ETX*End of TeXt.

ETB and *ETX* signal the receiving station that the line is to be turned around for an acknowledgment. *ITB*. on the other hand, can be used to denote the end of a record; it is followed by another record (typically, although not necessarily, starting with an *STX*). The number of *ITB* blocks allowed is implementation dependent; however, after a maximum of about four unacknowledged records, the fifth must be terminated by an *ETB* or an *ETX*. The *ETX* is a special case of *ETB*. indicating that no further frames will arrive (it is often used to signal that the end-of-file has been reached).

Frames consisting of printable bytes only (i.e., alphanumerics, carriage return, line feed, tab, and form feed) are known as nontransparent. Figure 6.56 shows the nontransparent frame format (note that, unless otherwise indicated, all frames are sent from left to right).

A 16-bit cyclic redundancy check (CRC-16) is calculated on all the data bytes and is transmitted after the end-of-frame character. The CRC covers only the data from the start-of-frame character (*STX* or *SOH*) to the end-of-frame character (*ITB*, *ETB*. or *ETX*). A separate CRC is calculated for each frame, including *ITB* frames. The CRC associated with an *ETB* or *ETX* following a series of *ITB* blocks is generated only on the last block.

The Bisync protocol allows a limited number of control characters to be

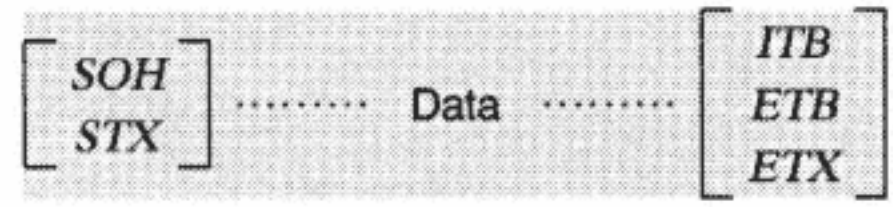

Figure 6.56
The nontransparent frame format.

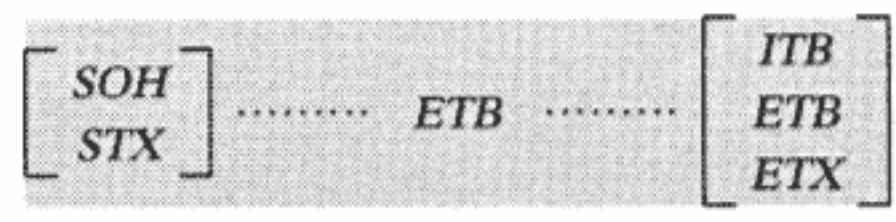

Figure 6.57
Nontransparent frames cannot include
certain control characters.

sent in nontransparent frames. This leads to another problem: how messages, such as executable images or binary files, consisting of nonprintable control characters, can be transmitted. For example, in Figure 6.57 part of the data stream consists of the control character *ETB*. The receiver cannot distinguish between the *ETB* indicating the end-of-message *ETB* and the data *ETB*. This means that the receiver will stop receiving after the first *ETB* (the *data ETB*!).

The solution to this problem is to have the transmitting DTE inform the receiving DTE that nonprintable characters are included in the transmission. A communication that permits the inclusion of control characters as data within the frame is known as a transparent transmission. A transparent transmission allows *any* byte to be transmitted as data. As shown in Figure 6.58, control characters in a transparent frame are prefixed with the *DLE* character.

Should an *ETB* or virtually any other control character be sent as data, it will be accepted as such because it is *not* prefixed with a *DLE*.

Although prefixing framing control characters with *DLE* is intended to ensure that *any* character can be sent as data, problems arise if a message contains a *DLE* followed by an end-of-frame character (i.e., one of *ITB*, *ETB*. or *ETX*). As shown in Figure 6.59, the receiving DTE cannot distinguish between *DLE* followed by *ETX* sent as data or sent as an end-of-frame sequence.

Since the *DLE* and the character that follows it are both data, neither can be left out of the transmission. The solution, shown in Figure 6.60, is to

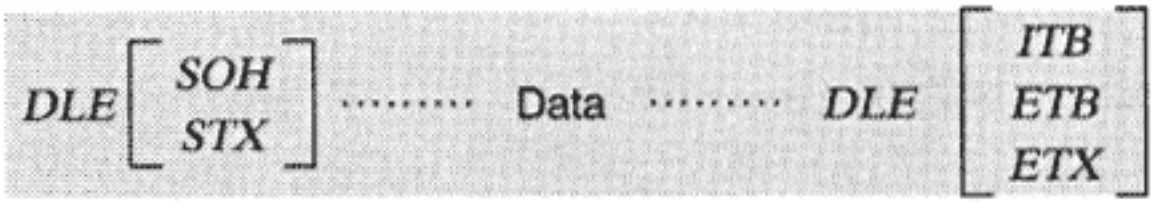

Figure 6.58
The transparent frame format.

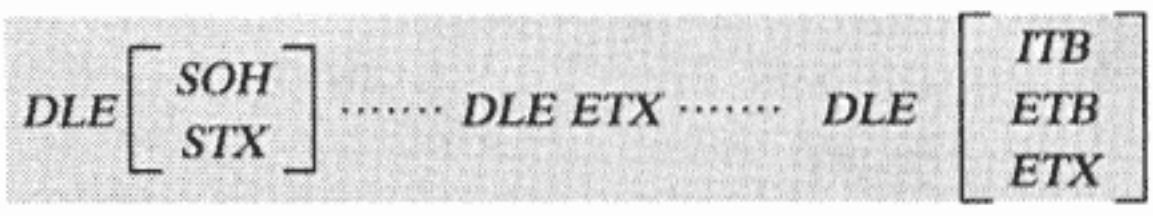

Figure 6.59
The *DLE-ETX* byte sequence cannot be sent in a frame.

Figure 6.60
All data *DLE*s must be prefixed by *DLE*.

prefix any *DLE* found in a message by another *DLE* and then to transmit both *DLE*s.

When the receiving DTE receives a *DLE* in the data stream, it enters a special state and waits for a *DLE* (taken as data) or a control character (typically signaling the end-of-frame).

Synchronization

Synchronization of Bisync DTEs is achieved through a series of synchronization characters (*SYN*) prefixed to the start of every frame. The *SYN* characters synchronize the transmitting and receiving clocks, thereby allowing the transmission of "pure" data. They can be placed in a frame to ensure that the receiver's clock maintains synchronization (there is yet another exception: if the frame is transparent, the *SYN* character must be prefixed by a *DLE*).

6.5.2
Bit-Oriented Synchronous Communications

Byte-oriented synchronous communications, such as those described above for Bisync, are still in widespread use, but they suffer from a number of drawbacks. These include the following:

Two frame structures are required: one for transparent data and the other for nontransparent data. In fact, there are numerous exceptions when dealing with the two frame structures, suggesting that Bisync is actually two protocols rather than one.

Using a *DLE* to signal that the next byte should be treated as data is an additional overhead in terms of both processing (adding and removing the *DLE*) and throughput (the transmission of 16 bits of information to represent 8 bits).

Both of these problems can be overcome if *bit* -rather than byte-oriented communications are used. A good example of a bit-oriented synchronous protocol is IBM's Synchronous Data Link Control, or SDLC.

In SDLC, all frames (containing information or acknowledgments) have the same structure, shown in Figure 6.61.

The fields in the SDLC frame are defined as follows:

Figure 6.61
The fields of an SDLC frame.

Flag. The Flag byte delimits the frame and is used to synchronize the two stations. Both the leading and trailing Flag bytes have the same pattern: `01111110`. SDLC allows multiple frames to be sent with a single Flag byte separating them.

Address. The Address field denotes the intended destination of the message (if the frame is sent by the Primary station) or the address of the station sending the frame (if the frame is sent by the Secondary station). This field is intended to allow several secondary stations to share a single communication channel, controlled by a single primary station through the use of polling.

Control. The Control field is one byte long and defines the function of the frame. A frame supports one of three formats (defining the frame's function):

Information Transfer Format signals the receiving station that the frame is an Information frame and contains one or more bytes of data. The format of the information transfer Control field is as follows:

Ns is the send count, which is the sequence number of this (information) frame. It is applicable only to Information frames. Nr is the receive count, indicating the sequence number of the next expected Information frame from the station to which this frame is destined.

Supervisory Format is intended to assist in the orderly flow of Information frames between stations. The Supervisory frame Control byte has the following format:

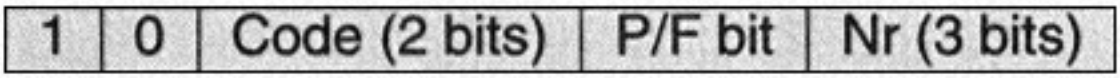

Code is one of the following:

Ready-to-Receive(00), or RR, is an acknowledgment; it signals that the next expected frame is Nr.

Receive-Not-Ready(01), or RNR, is both an acknowledgment (the next expected frame is Nr) and an indication to the transmitting station that there is a temporary problem at the receiving station.

Reject(10), or REJ, is a negative acknowledgment, requesting the transmitter to start transmission from frame Nr. Nr is the receive count, indicating the sequence number of the next expected information frame.

Unnumbered Format is used to convey information specific to the function of the communication channel connecting the various stations. For example, the Primary station can initialize Secondary stations through the use of Unnumbered frames, and stations can report their status in an Unnumbered frame. The format of an Unnumbered Control frame field is as follows:

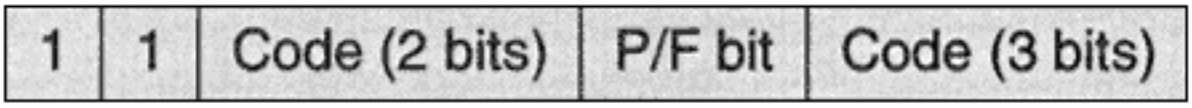

Code values include

UI (*Unnumbered Information*). The frame contains data in the Information field.

DISC (*Disconnect*). The Primary station is disconnecting the Secondary station specified in the Address field.

UA (*Unnumbered acknowledgment*). An affirmative response to a command such as DISC.

The P/F (Poll/Final) bit is common to all Control field formats; its interpretation is given in Table 6.5.

Table 6.5 The interpretation of the P/F bit.

P/F bit	Transmitter	Interpretation
Set	Primary (to Secondary)	Secondary is to start transmission.
Set	Secondary (to Primary)	Secondary has finished transmission.
Clear	Either	A transmission is in progress.

Information. Information contains data being sent from one station to another. If the frame is unnumbered, the Information field may be present, carrying information specific to the SDLC protocol such as a test pattern or station identification data. Information frames carrying data for applications on the destination DTE use the Information field; there is no specified length of the Information field, but it must be a multiple of 8 bits. Supervisory frames do not have an Information field.

FCS. The FCS is the Frame Check Sequence (another term for CRC), which is calculated using CRC-CCITT and covers the entire frame, excluding the Flag bytes. The FCS is generated by the transmitting station and placed after the Control or Information field (the Information field is optional in some frames); the receiving station calculates its own version of the FCS until the trailing Flag pattern is detected. The result of the receiving station's calculation should be `0xF0B8`.

SDLC uses a sliding window protocol with a total of eight possible sequence numbers; the window size is application dependent, with the maxi-

mum size being seven. The section on the go-back-N protocol describes the basic functions of SDLC.

Transparency

To be truly useful, a protocol must be able to support transparent data, thereby allowing any bit pattern to be transmitted. In byte-oriented protocols such as Bisync, a special escape byte (*DLE*) allows the receiving DTE to distinguish between control bytes sent as data and those used for control sequences. As well, the escape byte requires a set of special rules to handle its own inclusion.

In bit-oriented protocols such as SDLC, the problem of transparency still remains, since the sequence `01111110` delimits the end-of-frame, meaning that a byte sent with the value `01111110` will cause the receiving DTE to stop receiving. The solution is to use a technique known as bit stuffing, in which the transmitting hardware inserts a `0` bit after any sequence of five `1`s (except when the end-of-frame sequence is sent).

The receiving hardware removes the extra bit using the following algorithm:

1. If a sequence of five `1`s has been received (i.e., `11111`), wait for the next bit.

2. If the bit is a `0`, discard it and continue assembling the stream of bits into bytes.

3. If the bit is a `1`, keep itthis should be the end-of-frame sequence.

4. Reset the bit-stuffing sequence count and resume the count when the next `1` is received.

The following examples illustrate how bit stuffing works (0 bits inserted are enclosed as [0] and the information received has the inserted bit removed):

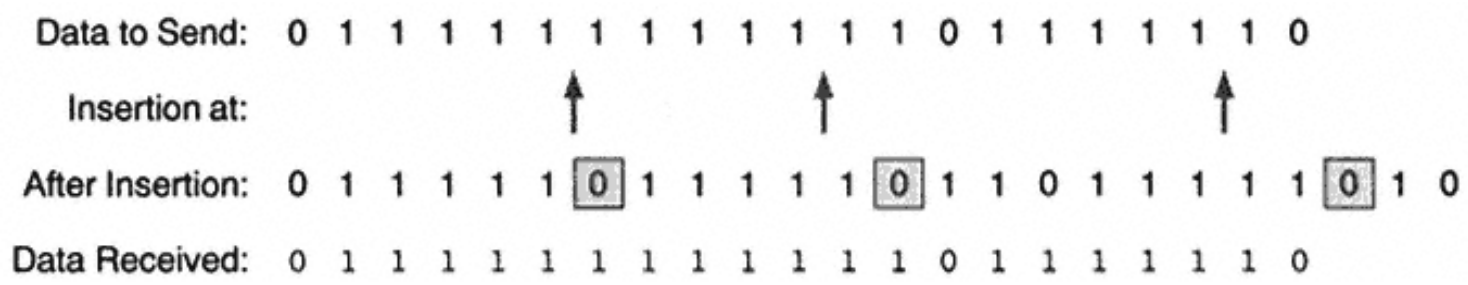

6.6
Synchronous Communications Hardware

Synchronous communications, like asynchronous communications, require specialized hardware to support the transmission and reception of data. Not surprisingly, there are several types of synchronous hardware available: at a minimum, it is either byte-or bit-oriented. The exact functions of the hardware vary from manufacturer to manufacturer; however, it is not surprising to find features such as synchronization detection and CRC generation and checking on the hardware itself.

6.6.1
Byte-Oriented Hardware

Many manufacturers of byte-oriented synchronous hardware make the hardware support asynchronous communications as well. Hardware that supports both types of communication is known as a USART or Universal Synchronous/ Asynchronous Receiver/Transmitter. The structure of the USART, shown in Figure 6.62, is similar to that of the UART.

The Intel 8251A is used on the IBM PC to support the Bisync protocol. Functionally, the 8251A USART is very similar to the 8250 UART in that the baud rate, character length, stop bits, and parity can be selected under program control. The program can also choose asynchronous or synchronous operation. (The 8251A supplied with the PC operates in synchronous mode only.)

When transmitting a Bisync frame, the software is required to format the entire frame (including the *SYN* characters and the CRC), which is transmitted one byte at a time. The software is expected to supply data to the 8251A at such a rate that the transmit buffer is never empty. Should the transmit buffer become empty, the 8251A starts to idle with *SYN* characters. This can clearly lead to problems in transparent mode, since *SYN* characters sent as timing characters are to be prefixed with *DLE*.

If a Bisync station is to receive a frame, the 8251A must be put into *hunt* mode, scanning the channel for one or more SYN characters. The software can also specify whether the search is for one or two *SYN* characters. The bytes are returned to the processor as they are received.

6.6.2
Bit-Oriented Hardware

Intel also manufactures the Intel 8273, the PC's programmable SDLC communications adapter. The 8273 performs a number of tasks specific to SDLC, thereby reducing processor overheads. For example, if a frame is to be transmitted, the 8273 transmits the Flag byte followed by the Address and Control

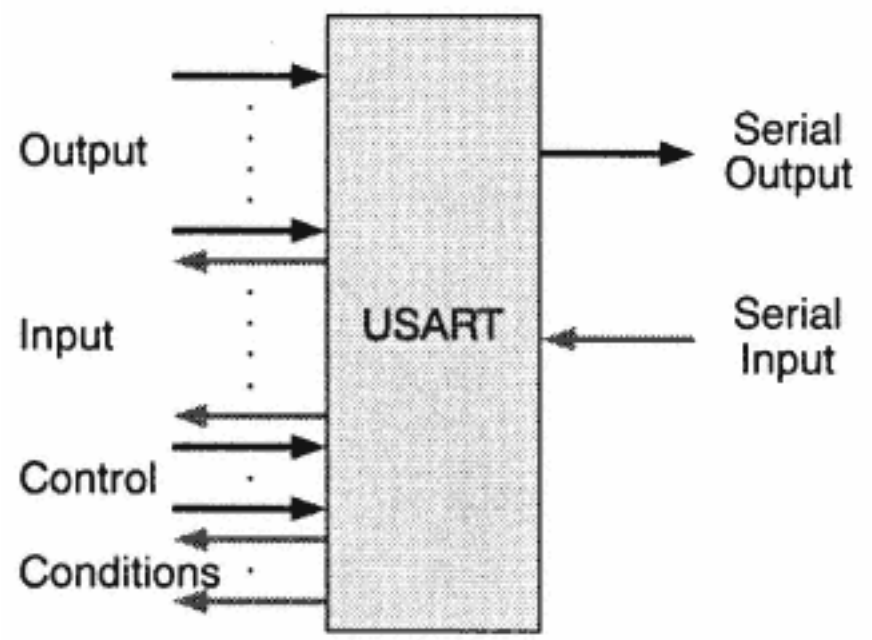

Figure 6.62
The USART.

fields. If an Information field is to be transmitted, it is supplied to the 8273 using direct memory access (DMA), thereby reducing processor overheads. Lastly, the FCS and final Flag byte are sent by the 8273. Incoming frames are treated in a similar fashion, with the Address and Control fields written to hardware registers and the Information field transferred to memory without processor intervention. The FCS is checked as well. The processor is interrupted upon reception of a frame.

6.7
Commkit Synchronous Software

Few PCs are supplied with synchronous communication hardware because of the costs involved. Thus, Commkit does not support software to control synchronous hardware, but it can *emulate* byte-oriented synchronous protocols on asynchronous hardware by using the same frame delimiters, error checking, and recovery procedures, and other rules associated with the protocol. The remainder of this chapter examines Commkit's half-duplex stop-and-wait protocol, which functions as a synchronous protocol.

6.7.1
Hardware Requirements

The stop-and-wait implementation is point-to-point and requires, at a minimum, two PCs. Either of the serial ports (1 or 2) can be interconnected. A better appreciation of the protocol can be obtained if the line analyzer is used. Figure 6.63 shows the configuration for monitoring the traffic between two synchronous stations.

6.7.2
The Problem

The problem has two parts: a stop-and-wait protocol is to be designed to support one-directional data transfer, and a high-level utility is to take a "block" of up to 10 lines of data and transfer it to the remote PC.

Figure 6.63
Hardware configuration for Commkit's "synchronous" software.

6.7.3
Design

A Stop-and-Wait Protocol

The stop-and-wait protocol will allow a process on one PC to send frames of up to 80 printable (i.e., nontransparent) characters to a process on another PC. The protocol is to be typically stop-and-wait, requiring the transmitting DTE to send a frame (with an embedded sequence number), to which the receiving DTE will reply with an acknowledgment: an *ACK* if the frame is received correctly or a *NAK* if an error is detected. If the frame is positively acknowledged, the transmitting DTE will send another frame; a negative acknowledgment will cause the transmitting DTE to retransmit the old frame. A frame received with the wrong sequence number is discarded but positively acknowledged. If no response is received within one second of the transmission, the frame is again sent. A possible sequence is shown in Figure 6.64.

The protocol uses Bisync control characters to delimit each data frame. The frame structure is shown in Figure 6.65.

The sequence number seq alternates between 0 and 1, while the Data consists of from 0 to 80 bytes. Any control character (other than *ETX*) can be sent as data. The 16-bit CRC, generated using Commkit's CRC-16 utility and

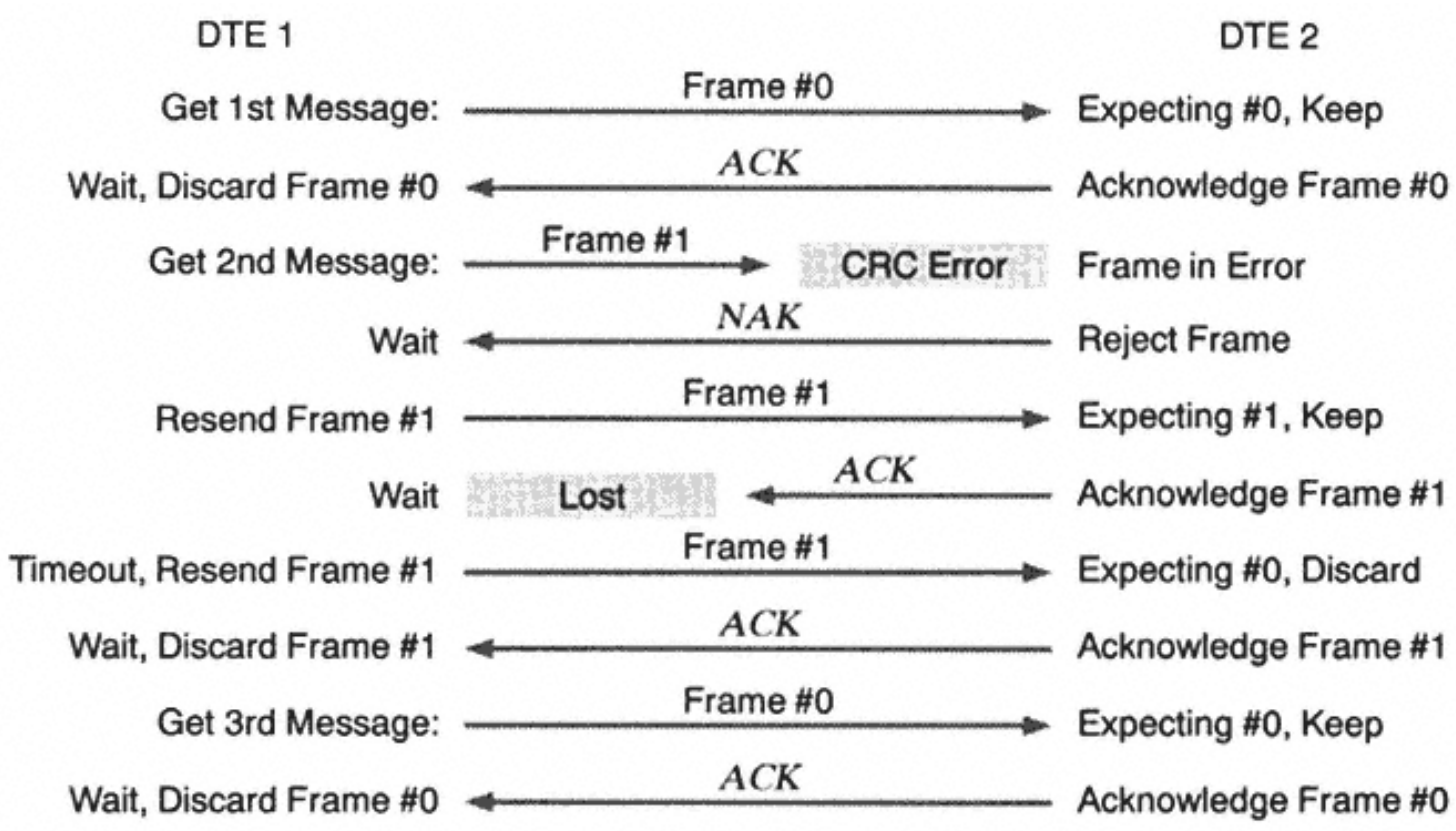

Figure 6.64
The Commkit stop-and-wait protocol.

Figure 6.65
The Commkit frame.

sent as two 8-bit bytes (CRC) includes *both* Data and the sequence number seq. The data to be transmitted is supplied by a process; once the transmission has completed, a completion code is to be returned to the process indicating that the next line of data can be transmitted.

The acknowledgment characters (*ACK* and *NAK*) are the ASCII positive and negative acknowledgment characters, respectively.

Block Transfer

There are *two* foreground processes:

On the transmitting PC, the foreground process will accept up to 10 lines of text from the user and, when signaled, send it to the stop-and-wait software for transmission to the receiving DTE. When all the lines on the transmitting station have been sent, the screen is cleared, ready for further input. Data is forwarded if one of two conditions is met: first, whenever ten lines of text are entered; second, a CTRL-Z is pressed, which will cause all messages (at least one and less than ten) to be sent.

On the receiving PC, when a line is received correctly by the stop-and-wait software, it will be forwarded to the local foreground process for display. Whenever the number of lines on the screen exceeds 10, the foreground process will clear the screen and proceed to display the incoming lines.

Both stations can act as the transmitter or receiver, but they are not to change roles. The transmitting station can only transmit blocks of data, while the receiving station can only receive them.

6.7.4
Implementation

The stop-and-wait software is implemented as two separate files, one containing the "low-level" routines supporting message transfer, and the other a set of "high-level" routines, responsible for all message and keyboard display (the division of the software into these two levels will be used in subsequent chapters to demonstrate other issues in data communications). The two stopand-wait files are

`s&wlow.c`. The procedures in `s&wlow.c` are responsible for the low-level transmission and reception of messages across the channel. Messages are read from the queue indicated in `handler_id` and sent a byte at a time. The port is specified by the user.
A station receiving a frame responds with an acknowledgment (*ACK* or *NAK*), extracts the message from the frame, and sends it to the process associated with queue `APPLICATION`.

`s&wex.c` . `S&wex.c` contains the high-level software responsible for displaying messages received from queue `APPLICATION`. All keyboard

input is also sent to APPLICATION for display and formatting into messages. Once 10 lines or CTRL-Z is typed, the messages are sent to queue handler_id for transmission.

This problem can be divided into four parts: high-level transmission, high-level reception, low-level transmission, and low-level reception.

High-Level Transmission

The high-level transmission software, do_1c1_msg(), is in one of two states:

In state READING, characters are accepted from the keyboard and stored in the structure msg_to_go until a CTRL-Z or end-of-buffer is detected. Structure msg_to_go is an array of 10 elements (one for each line):

```
struct
{
int size;
char data[81];
} msg_to_go [10];
```

As each character is accepted, it is displayed on the screen at location line, column. The character is stored in msg_to_go.data, and column is assigned to msg_to_go.size.

When msg_to_go is filled or the user signals CTRL-Z, do_1c1_msg() begins its writing sequence by calling write_message_number(), which sends the first line stored in msg_to_go . data to the low-level transmission routine. All messages are sent to one of the serial ports (SP1IH or SP2IH) indicated by handler_id. The state then changes to WRITING. The variable current_msg contains the number of the next line to be transmitted.

A CTRL-C from the keyboard will abort the process, returning control to MS-DOS.

The WRITING state is entered whenever the low-level transmission routine has completed sending a message to the remote PC. If there is another line of data to be sent (indicated by current_msg being less than line), write_message_number() is called with current_msg.

Control remains in the WRITING state until all the lines of data are transmitted, at which point the state changes to READING.

High-Level Reception

Frames received by the low-level reception software are forwarded to do_rmt_msg(), the foreground process, through queue APPLICATION for display. The lines are displayed as received starting at the top of the screen; after the tenth line, the screen is cleared and display starts again from the top line.

In both the high-level reception and transmission software, there are two common routines: `diagnostic()`, which writes a diagnostic message on the last line of the screen; and `check_clock()`, which erases the last line of the screen after about five seconds of display have elapsed.

Low-Level Transmission

All entry to the transmission software takes place through `low_level()`. Interrupts from either serial port or the clock handler are eventually routed to `xmit_protocol()`, which is responsible for the transmission of the frame. `Xmit_protocol()` is implemented as a state machine (see Figure 6.66) consisting of the following states; the value of the current state is stored in `s_and_w.state` (all data structures are defined in `s&w.h`):

`AWAIT_MSG`. State `AWAIT_MSG` is entered initially and after a message has been successfully transmitted. If a message is available for transmis-

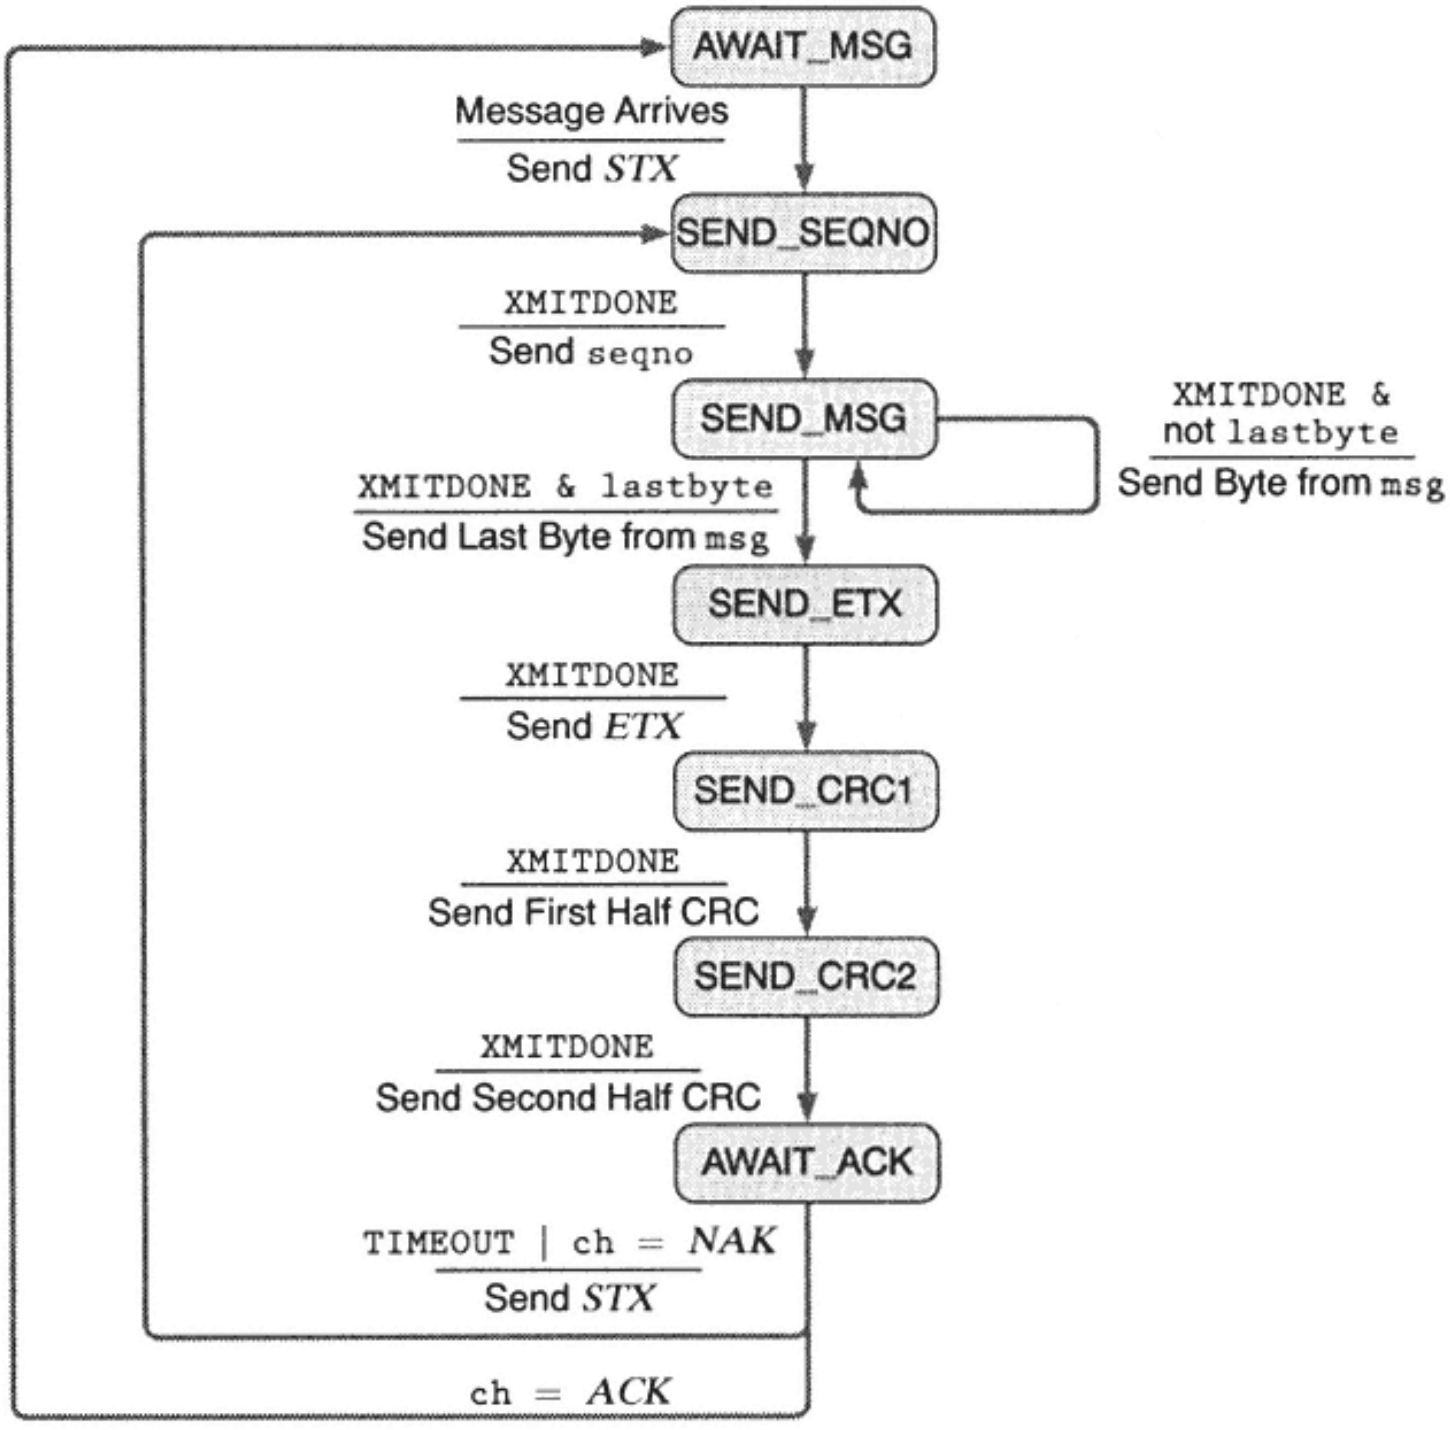

Figure 6.66
Stop-and-wait protocol: low-level transmission state diagram.

sion, `start_transmission()` is called and an *STX* is sent. The state changes to `SEND_SEQNO`.

`SEND_SEQNO`. As soon as the serial port signals that the *STX* has been sent, control passes to state `SEND_SEQNO`. The sequence number, `s_and_w. seq_no`, is transmitted and included in the CRC. The state changes to `SEND_MSG`.

`SEND_MSG`. The message bytes, stored in `s_and_w.msg`, are transmitted as each `XMITDONE` indication is received. Each byte is included in the CRC. When end-of-message is detected and the last byte has been sent, the state changes to `SEND_ETX`.

`SEND_ETX`. `SEND_ETX` is entered after the last byte of the message has been sent. An *ETX* is transmitted, and the state changes to `SEND_CRC1`.

`SEND_CRC1`. The first half of the CRC is sent; the state changes to `SEND_CRC2`.

`SEND_CRC2`. The second half of the CRC is sent; the state changes to `AWAIT_ACK`. The timer `s_and_w.ticks` is cleared.

`AWAIT_ACK`. The `AWAIT_ACK` state is entered after a message has been sent; the transmission state machine is awaiting an acknowledgment from the remote station. There are three possible outcomes:

No response. If no response is received after one second, the frame is resent by calling `start_transmission()`. The state changes to `SEND_SEQNO`.

Response is *NAK*. A negative acknowledgment results in the frame being retransmitted by a call to `start_transmission()`. The state changes to `SEND_SEQNO`.

Response is *ACK*. A positive acknowledgment means that the remote station has received and accepted the frame just sent. At this point, the sequence number is updated and the high-level software informed of the successful transmission. The state changes to `AWAIT_MSG`.

Any other responses are ignored.

Low-Level Reception

Entry to the low-level-reception software takes place through `low_level()` whenever an interrupt is received from either the serial port or the clock handler. `Recv_protocol()` handles the reception and acknowledgment of any incoming frames through the use of a state machine, the state of which is indicated in `s_and_w.state`. The action of the state machine depends upon the current state and the condition for changing state (see Figure 6.67):

`WAIT_STX. WAIT_STX` is entered initially and after a message has been received. When an *STX* has been received, control passes to state `WAIT_SEQNO`.

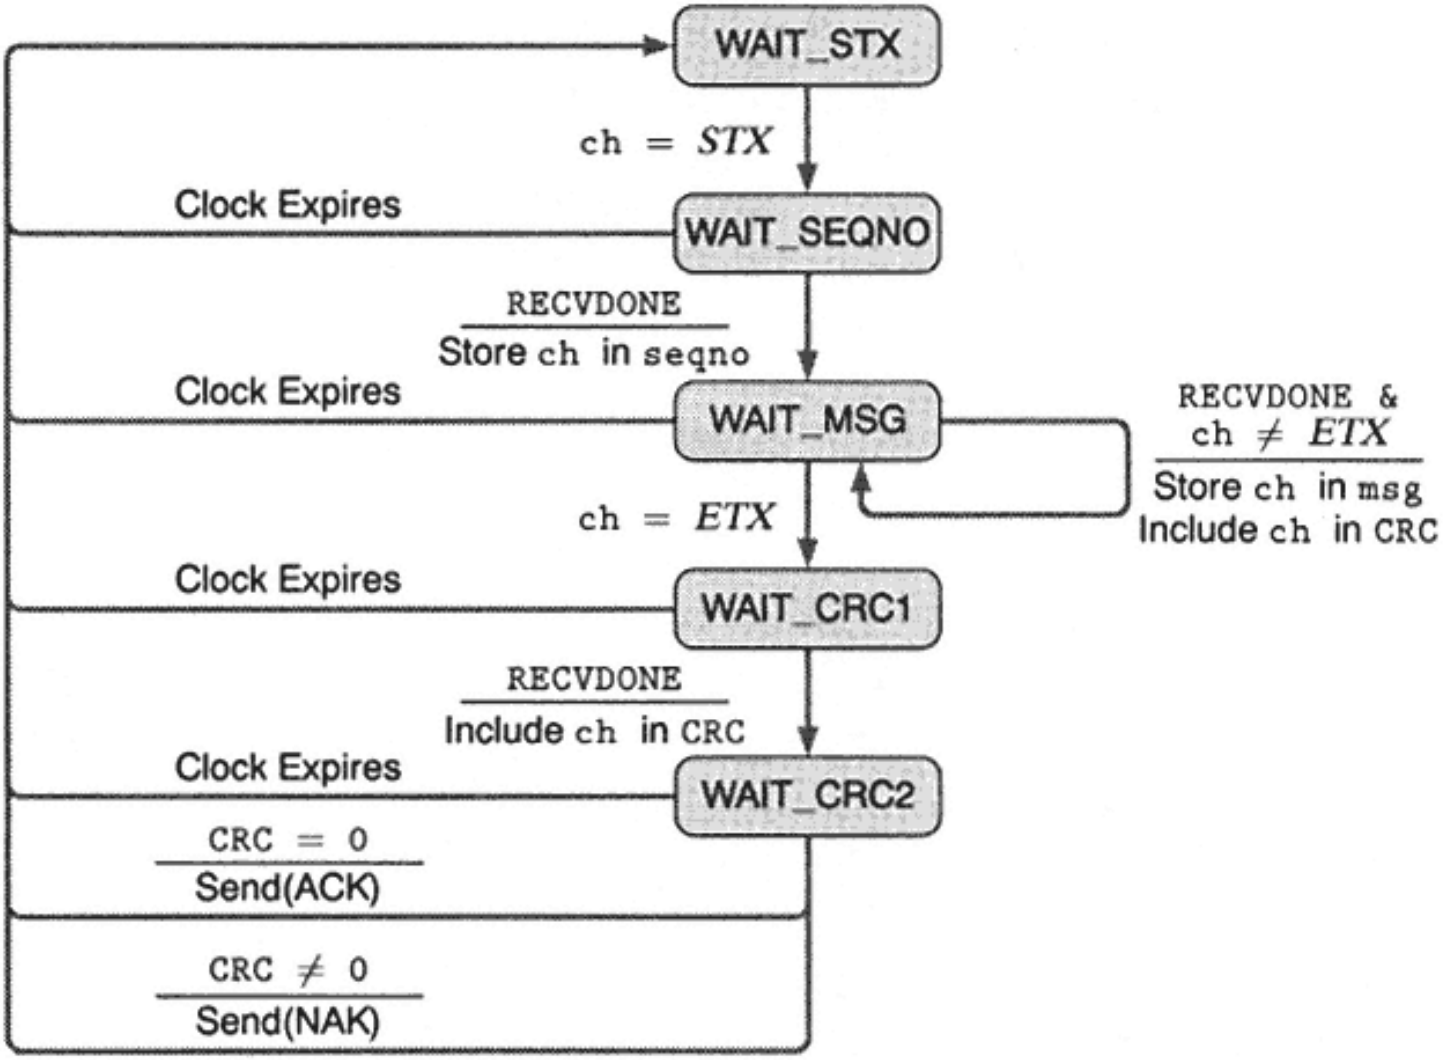

Figure 6.67
Stop-and-wait protocol: low-level reception state diagram.

`WAIT_SEQNO`. The byte following the *STX* is taken to be the sequence number of the frame and is stored in `s_and_w. recv_no`. The sequence number is included in the CRC. The state then changes to `WAIT_MSG`.

`WAIT_MSG`. The bytes following the sequence number are stored in the array `s_and_w.msg` and included in the CRC. When an *ETX* is detected, control passes to state `WAIT_CRC1`.

`WAIT_CRC1`. The byte following the *ETX* is taken as the first half of the CRC. This byte is included in the CRC being calculated for the frame. The state then changes to `WAIT_CRC2`.

`WAIT_CRC2`. This byte is taken to be the second half of the CRC and included in the CRC. `Process_message()` is then called, which examines `s_and_w.crc`. If the value of the CRC is 0, the frame is acknowledged; if the sequence number is the one expected (indicated by comparing `s_and_w. recv_no` with `s_and_w.seq_no`), the message is forwarded to queue `APPLICATION` and the sequence number is incremented.

In each state (except `WAIT_STX`), there is a timed sanity check; should the transmitting station fail to send a byte within one second, a diagnostic message is displayed on line 0 and control passes back to state `WAIT_STX`.

The state of the low-level software (either transmitting or receiving) is indicated in `s_and_w.protocol`: `XMIT` or `RECV`. All keyboard messages are sent to the foreground process.

6.7.5
Compiling and Running the Stop-and-Wait Software

The stop-and-wait software is also supplied on the Commkit diskette as the executable file s&wex. exe. A new executable can be created with the make utility:

```
C:\> make s&wex.exe
```

`S&wex.exe` is created from `s&wex.c`, `s&wlow.c`, `commkit.obj`, `crc.obj`, and `srcalls.obj`.

Once created, `s&wex.exe` can be executed by typing the file name followed by the desired line speed, the port in question (1 or 2), and an indication of the function of the PC, either transmitting (*X*) or receiving (*R*the default). For example, to run a PC as a transmitter at 300 bits-per-second through port 2, type:

```
C:\> s&wex 300 2 X
```

Both stations (transmitter and receiver) cause the screen to be cleared upon initialization. If the above format is not followed or the line speed is not recognized, control returns to MS-DOS and a diagnostic is issued. Data is entered at the PC designated as the transmitter. Either station can be terminated by typing CTRL-C.

As data is entered on the transmitting station, it is echoed on the screen. A new line is obtained after each carriage return, or the user attempts to enter a character on a full line. The contents of the screen are forwarded a line at a time until 10 lines have been entered and the user attempts to move to the 11th or types CTRL-Z.

The version of `s&wex.exe` supplied on the distribution diskette will automatically start a new line after the 61st character is entered. This permits the transmitting station to send the same message repeatedly to the receiving station without user interactiona useful feature when testing the software.

Each line, as it is received, is displayed on the receiving PC, starting on the second line down from the top of the screen. Any errors detected by the receiving PC (such as bad sequence number or invalid CRC) are displayed on the top line.

Once the transmission has completed, the transmitting PC's screen is cleared and the user can enter more data. The receiving PC's screen is cleared before the first line of any incoming data is displayed.

6.8
Summary

This chapter examined synchronous communications and their usesnotably, the transfer of large volumes of data with the minimum of overheads. Many of

the issues involved in synchronous communications were discussed, including the following:

Synchronous communications involve the transmission of "pure" data (i.e., no start and stop bits); to achieve this, special frames are needed with some form of delimiter.

The receiving DTE can verify that the transmitted frame has been received correctly through the use of a series of one or more check bytes that are generated by the transmitter. Techniques such as the longitudinal redundancy check (LRC) can be used, but they offer limited checking capabilities on the data within the frame. More accurate frame checking is possible using cyclic redundancy checks, or CRCs. A CRC treats the entire message as a long dividend and divides the message by a CRC generator; the remainder is transmitted as part of the frame. A number of CRC standards exist; CRCs can be generated in hardware and software.

Once the frame has been received, the receiving station needs a mechanism whereby the transmitting station can be informed of the status of the message (was it received correctly or was the CRC in error?). Such a mechanism is the acknowledgment, in which the station receiving the frame replies to the transmission either positively or negatively.

In environments where acknowledgments are not possible, forward error recovery allows the receiver to correct errors detected in the frame.

Two common classifications of protocol have been developed for use with synchronous communications: stop-and-wait and go-back-N. Stop-and-wait protocols involve a transmitting station sending a frame and then waiting for an acknowledgment, at which point the next frame can be sent. Go-back-N protocols permit potentially higher throughput by the use of sliding windows, in which the transmitter can send a number of frames before an acknowledgment is required. Go-back-N protocols also allow for the bidirectional transfer of information by piggybacking acknowledgments on data frames.

The chapter also demonstrated that synchronous communications can be emulated, in part, by asynchronous communications. The Commkit stop-and-wait protocol is one example.

6.9 Exercises

1. (A) Show that synchronous communications offer better throughput than asynchronous communications.

2. (A) If synchronous communications are so much more efficient that asynchronous communications, does it make sense to replace all asynchronous terminals with synchronous terminals? Consider the overheads.

3. (A) Compare the overheads involved in the calculation of a message's CRC using the software shift-register implementation against the table lookup method.

4. (A) Show, by means of illustrations or examples, why throughput can be improved in the stop-and-wait protocol by using a *NAK* rather than a timeout.

5. (A) Show, by means of illustrations or examples, how *NAK*s can improve throughput for selective retransmission.

6. (A) Show, by means of illustrations or examples, how a sliding window protocol can be used to support selective retransmission.

7. (A) Show the encoding of the following nibbles using the Hamming code: `0000`, `1001`, `0110`, and `1010`. Verify the results by decoding the encoded data.

8. (A) Indicate which of the following 7-bit Hamming frames are correct; correct those that are incorrect: `0000010`, `1111011`, `1111111`, and `1000111`.

9. (A) In Bisync, why does a timing *SYN* require a *DLE* prefix when transmitting transparently? What happens if the *DLE* is omitted?

10. (A) How does SDLC distinguish between its three frame formats?

11. (A) Using the state machine descriptions for the stop-and-wait protocol, show how the protocol recovers from the error shown in Figure 6.68.

12. (B) Run `s&wex.exe` between two PCs connected by the line analyzer. Identify the different fields in the frame. What are the values of the sequence numbers?

13. (B) Run `s&wex.exe` from one PC connected through the line analyzer. What happens when several lines of message are sent? How many messages are sent? Why?

On the line analyzer, the frame delimiters appear as ☺ (*STX*) and ♡ (*ETX*), while the acknowledgment is displayed as ♠ (*ACK*). Note that the CRC bytes follow the *ETX*.

14. (B) Run `s&wex.exe` between two PCs connected by the line analyzer. Fill a screen (i.e., 10 lines) with the same character sequence; transmit the messages. Explain the differences in the CRC values displayed on the analyzer.

15. (C) Modify the `s&wex.exe` software so that CTRL-C does not terminate the transmitting PC. Next, run `s&wex.exe` between two PCs connected by the line analyzer. Insert several *ETX* characters as data into the mes-

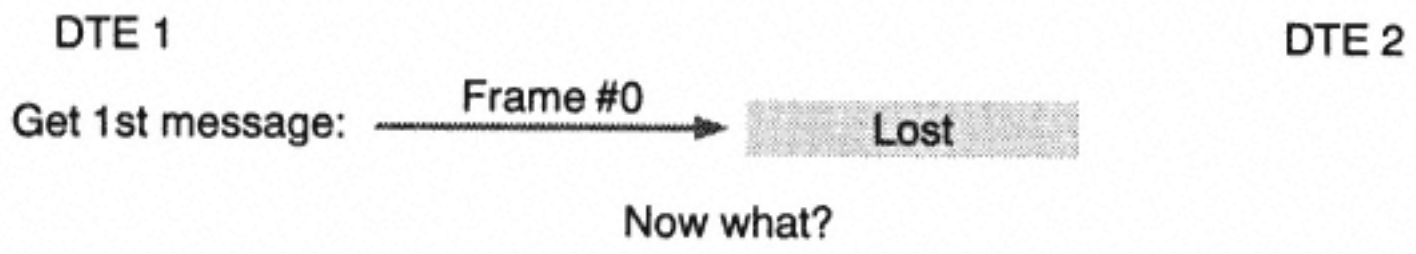

Figure 6.68
How does the stop-and-wait protocol recover?

sages (type CTRL-C; it echos as a heart). What happens when the message is transmitted?

What modifications are necessary to the stop-and-wait software to allow the transmission of an *ETX* as data? Make the modifications and retest the software.

16. (C) The stop-and-wait protocol supplied with Commkit has a problem: if the receiving DTE misses the *ETX* and the transmitting DTE retransmits the frame before the receiving DTE can abort the first reception, a very long message will be received, possibly causing various areas of memory to be overwritten. Show where the problem occurs in the state diagram; is it a transmission or a reception problem (i.e., which fails, the transmitter or the receiver?)? Modify the existing state diagram and the software to handle the error described here.

17. (C) The example stop-and-wait protocol has another problem: if a frame is continuously *NAK*ed, the transmitter will never stop transmitting and the receiver will never stop receiving. A constant `SEND_ABORT` has been defined in `s&w.h`, but isn't used. Modify the stop-and-wait software so that after five retransmissions of the same packet, the transmission will be aborted. Develop a way for this to be tested. What changes are needed in `s&wlow.c`?

18. (C) Many test situations may require an error to be introduced into the transmitted data. For example, in order to check an implementation's error-handling routines, it is necessary to send messages with invalid CRC values. (It may be easier to have the line analyzer cause the errors than to rewrite the transmission software to generate errors!)

Modify the line analyzer software so that the user may enter data that is subsequently sent out a specified port. Use CTRL-J to enter the data to be sent and CTRL-V to send the data. The port out of which the data is to be sent can be specified in the first data byte entered.

7
File Transfer

7.1
Introduction

The examples illustrating both asynchronous and synchronous point-to-point communications have thus far been confined to short messages entered by the user at a PC. Not surprisingly, this is only one of many types of application that can use a communication facility. There are others, such as file transfer, involving the transfer of a file from one DTE to another. Protocols designed for large-volume transfer, such as those discussed in Chapter 6, are ideal for file transfer, since large numbers of bytes can be transferred reliably with the minimum of acknowledgment overheads.

Although the concept of file transfer is straightforward (a copy of a file is sent from one DTE to another), there are, as usual, a number of important issues that must be considered to ensure that the file arrives at the intended destination complete and in the form in which it left the source. For example:

Are the file attributes transferred? File attributes are information about the file, such as its name and extension, ownership, time and date of creation, size, and type (i.e., sequential, direct, or indexed). At a minimum, the file name and extension should be transferred. However, there are situations in which it is necessary to transfer all of the file's attributes. This may not be as simple as it seems. If an indexed file is to be transferred and the index refers to absolute disk block locations on the original disk, transferring the index may be a meaningless operation. Similarly, one type of file may not be supported on the destination DTE, requiring the file to be converted to a common file type. A case in point would be transferring an indexed file with its indexes to a DTE that supports only sequential access.

What will happen to the file on the destination DTE if the file transfer fails? File transfer can fail if one DTE (or both) or the communications channel ceases to function. If part of the file has already been written to the disk, should it be left on the disk or deleted?

Should an acknowledgment indicating that a message has been properly received by the remote DTE also be taken to mean that the message has been successfully written to the disk? In some situations, a message may be received correctly by the communications software, but may not be written correctly to the disk (the disk may fail, for example). *Additional* protocols may be needed to indicate that the message has been received *and* successfully written to the file on the disk.

This chapter presents an overview of how file systems function and how files can be transferred. The basic concepts associated with file transfer are examined in detail. A general-purpose file transfer application using Commkit's stop-and-wait software developed in Chapter 6 is discussed and implemented.

7.2 File Systems

Central to any file transfer is the operating system's file system. The file system is responsible for the organization of and access to files stored on some external medium, typically a disk. There are many approaches to the development and implementation of a file system; fortunately, a number of operations are common to most file systems, some of which are discussed in this section.

7.2.1 Disks

Disks are a magnetic (or electronic) medium that can store information. Information on a disk is divided into collections of bytes, referred to as sectors. Sectors are organized into rings on the diska single ring is known as a track. Information on the disk is read from and written to by a read-write head moving from track to track as the disk rotates. If a disk has multiple surfaces, each surface has its own read-write head that moves across all tracks. A cylinder is defined as the tracks of a multiple-surface disk that can be accessed without moving the read-write heads.

In many file system implementations, sectors are grouped into blocks; disk access requests are expressed in terms of blocks rather than individual sectors.

7.2.2 Files

Files consist of a series of bytes often grouped into structures known as records and stored in one or more blocks on the disk. These blocks are maintained in a structure (usually a list structure, linked or sequential). The file system maintains a directory (also made up of a number of blocks), which contains the names of the files on the disk. The attributes associated with each file are kept on the disk as well, often as part of the directory or in a separate data structure. One of the attributes is a pointer to the file's first block.

7.2.3
File Access

At a minimum, a file system allows a process to access files for both reading and writing. Typical file system operations include

Creation. Initially, before any files are put on it, the disk contains an empty directory and a list of available blocks. Processes that require the storage of information can request the creation of a file. The file system adds the name of the file to the file directory along with any other necessary attributes. Once the file is created, the process may write information to it.

Opening. An existing file (one with a name in the directory) is accessed by the application requesting that the file be opened, usually for reading or writing. Many file systems will check the attributes associated with the file before performing the open request; for example, only processes with certain privileges may be allowed to open a file.

There are numerous special cases of file opening. For example, a file can be opened in "append" mode, allowing the application to write to the end of it and thereby preserve its original contents. Other file systems can open files for both reading and writing, permitting the updating of specific records in them. Finally, file creation is simply a special case of opening a file for writing from its beginning.

Closing. When a process has finished its file access, the file system usually expects the application to close the file. By the closing of the file, other processes can be granted access to it (assuming that there is a limit of one process per open file).

Reading. Files are opened for reading to allow processes access to the information within them. Exactly *how* the file is read depends upon the file type and the file system. For example, sequential access returns the next available record to the process, whereas direct access can return any record in the file. When a sequential-access file is read, mechanisms exist to signal the process that the end-of-file has been reached.

Writing. A file can also be opened for writing, thereby allowing the process to update existing information or add entirely new information to it. As the process adds information, the file system places it into blocks and, when the block is filled, writes the block to the disk.

Since the file system may allow many files to be opened simultaneously, each file opened is associated with a handle, or file descriptor. The file descriptor is used by the file system to distinguish between the various file requests from the different processes, and it is used by the process to distinguish between the different files it may have open.

Other commands associated with file manipulation can be made out of the above operations. For example, a file copy involves creating a new file and opening an old one. The old file is read, and each record is written to the file newly created. Once the copying has completed, both files are closed.

7.3
File Transfer Protocols

A file transfer protocol is a set of rules describing the steps required for a file transfer to take place. There is no single set of rules for file transfer protocols; for example, some are full-duplex, permitting simultaneous bidirectional file transfer and others are half-duplex, allowing one file to be sent at a time. Similarly, some file transfer protocols operate using a stop-and-wait protocol while others are implemented using go-back-N protocols.

In its simplest form, a file transfer protocol can be thought of as a file copy command operating between a pair of DTEs. Ideally, the protocol will be written so that it is independent of the underlying communication software, as shown in Figure 7.1.

By separating the file transfer software from the communication software, a number of benefits can be obtained, including

Portability. If the file transfer protocol makes no references to the communication protocol, changing the communication protocol should be transparent to it. For example, if one pair of DTEs uses a stop-and-wait communication protocol and another pair uses a go-back-N protocol, a truly portable file transfer protocol will operate atop either protocol (this assumes that both communication systems offer the same set of communication primitives to the file transfer software).

Testing and verification. Software testing and subsequent verification are simplified if the different layers are written independently of each other. Quite simply, the number of places in which an error must be searched for is greatly reduced if each part of the system is tested independently.

Given this, the underlying communications software and its implementation can affect the types of files that can be transferred. For example, if the communication software supports only 7-bit ASCII data and the file contains 8-bit binary data, file transfer may not be be possible.

Regardless of the implementation, a file transfer protocol describes three basic operations:

File identification. The existing file must be opened for reading on one DTE and the new file created on the other DTE. In both cases, a file name

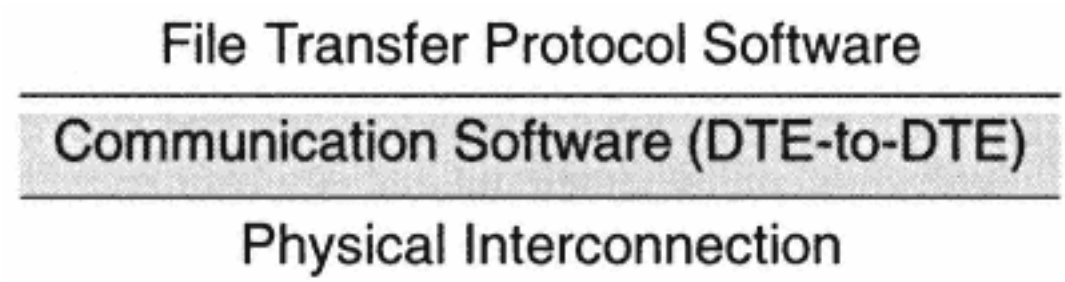

Figure 7.1
File transfer software should be independent
of the physical interconnection.

(and possibly a path indicating the storage unit and the directory) must be specified on both machines. Additionally, file attributes must be associated with the new file; these attributes can be those of the original file, or they can be the default attributes assigned by the file system on which the file is created.

The file transfer can be aborted at this point for any number of reasons (note that most of these points are implementation dependent):

The file to be copied does not exist, so the transfer cannot take place. (Some implementations will ask the user for the name of another file.)

The file name is already in use on the receiving DTE. A common solution is to rename the file to something else rather than to abort the transfer.

The user requesting the transfer does not have the necessary access rights to the file. Some files will be associated with certain privileges that dictate who can access them.

Where the file transfer is initiated is, once again, implementation dependent. The simplest approach is to have the transfer initiated on the DTE where the file exists and to have a user on the remote DTE specify the name and path of the file to be created. Ideally, a user on either DTE can request a transfer to (or from) the other DTE. If the transfer can be initiated from either DTE, regardless of the file's location, the file name must be passed between the two with an indication as to whether the file is to be created or opened for transfer.

Record transfer. Once the file has been opened for reading on one DTE (the source) and created for writing on the other DTE (the destination), its contents must be transferred. In their simplest forms, the algorithms for transferring the contents of a file can be summarized as follows:

The source algorithm:

1. Open the file for reading.
2. Read a record from the file.
3. Pass the record to the communication software for transmission.
4. Repeat steps 2 and 3 until end-of-file is detected.

The destination algorithm:

1. Create the file.
2. Wait for a record from the communication software.
3. Write the record to the newly created file.
4. Repeat steps 2 and 3 until end-of-file is indicated.

Although these are the basic steps in transferring the contents of the file, the algorithms are very often extended in light of the following:

It is not always advisable to allow the unchecked transfer of records as suggested in the source algorithm. If the communication software runs more slowly than the file transfer software, records can be lost if the communication software's queue overflows. One solution is to use a so-called handshake between the file transfer and communication

software: the file transfer software supplies another message only when permitted to do so by the communication software.

If the communication software is prone to losing messages, many file transfer protocols will implement end-to-end acknowledgments between their two halves. These acknowledgments are handled by the protocol independently of the communication software, meaning that there are potentially two sets of acknowledgments taking place: those in the communication software and those in the file transfer software.

Finally, it is not a good idea to assume that the receipt of an acknowledgment indicates that the message just sent was actually written to the remote file. It is possible that the record was received correctly but not written to the file; for example, the disk might have failed during a write operation.

End-of-file indication. As already suggested in Section 7.3, on transferring the contents of the file, an end-of-file indication must be sent to the destination. It is typically *not* written to the fileit is simply a signal to the file transfer software to close the file. The file system takes care of updating the end-of-file information.

7.4 Commkit File Transfer Example

Commkit is supplied with `s&wdisk.c`, a file transfer utility that can transfer files between PCs using the low-level stop-and-wait software `s&wlow.c . S&wd.isk.c` is a simple file transfer utility that operates in one of two modes:

Source. In source mode, `s&wdisk` reads the contents of a given file and supplies each record, one at a time, to `s&wlow` for transmission.

Destination. In destination mode, `s&wdisk` waits for records received by `s&wlow` and writes them to the newly created file.

7.4.1 Design

The file transfer utility, `s&wdisk`, is written independently of the `s&wlow` software but designed so that `s&wlow` can support the transfer without any modification. Accordingly, there are a number of restrictions placed upon the design of `s&wdisk` because of decisions taken in the original design of `s&wlow`:

A maximum of 80 bytes can be transferred.

Communications are half-duplex; the PC initiating the transmission is the only one that can send messages, and the receiving PC can send only acknowledgments.

In light of these restrictions, the file transfer can be initiated only from the PC that has the file to transfer. Additionally, end-to-end acknowledgments

between the file transfer software is impossible because of the half-duplex restriction. Finally, recall that `s&wlow` signals to the transmitting process that the last message sent has been acknowledged.

The following algorithm supports file transfer from the source PC using `s&wlow`:

1. Open the file for reading.

2. Read a record from the file.

3. If end-of-file is not detected:

    a. Pass the record to `s&wlow` for transmission.

    b. Wait for `s&wlow` to indicate that the acknowledgment was received.

4. Repeat steps 2 and 3 until end-of-file is detected.

5. Pass an end-of-file indication to `s&wlow` for transmission.

The following algorithm uses `s&wlow` for reception on the destination PC:

1. Create the file.

2. Wait for a record from `s&wlow`.

3. If the record does not indicate end-of-file, write the record to the newly created file.

4. Repeat steps 2 and 3 until end-of-file is indicated.

5. Close the file.

In addition to the above algorithms, a message structure is required to permit the transfer of the individual records. As shown in Table 7.1, the data structure requires two fields, one specifying an option and the other specifying the data associated with the option.

7.4.2 Implementation

The two algorithms described in Section 7.4.1 can be implemented directly in Turbo C using the UNIX file access functions `open()`, `close()`, `read()`, and `write()` (these functions and their associated options are defined in the include files `fcntl.h` and `sys\stat.h`). Communications between the foreground process and `s&wlow` are supported by `send()` and `recv()`.

To simplify the design, the file name must be specified on both the source

Table 7.1 The file transfer protocol message structure.

Option Field	Data Field
Data	File record
End-of-file	Empty

and destination PCs. The file name is not transferred (implementation of this is left as an exercise).

Two mutually exclusive foreground processes implement the source and destination algorithms (which one is called depends upon whether the PC is the source or the destination of the file):

`read_and_send()`. The `read_and_send()` procedure opens the file specified by the user. Each record is read, `REC_LEN` bytes at a time, into the structure `message.data`. If end-of-file is detected, `message.option` is set to `EOF`; otherwise, it is set to `MORE_TO_COME`.

Message is forwarded to the serial port specified in `handler_id` for transmission by `s&wlow. Read_and_send ()` then waits for an indication that the transmission has completed or until a CTRL-C from the keyboard is detected. Note that `read_and_send()` does *not* terminate as soon as end-of-file is detected; this is to allow `s&wlow` the time needed to transfer the `EOF` message. Once the transfer has completed, the file is closed.

`receive_and_write(). Receive_and_write()` creates a file of the name specified by the user. The attributes associated with the file are `S_IREAD` and `S_IWRITE` (indicating that the file can be read and written).

Once the file is opened, `receive_and_write()` accepts each message (containing the `option` and `data`); if the option is `MORE_TO_COME`, the `data` is written to the newly created file. When `message.option` indicates `EOF`, the file is closed and `receive_and_write()` terminates.

The overall flow of messages between the source and destination PCs is shown in Figure 7.2, where

1. `S&wdisk` uses the `send()` primitive to pass `message` to `s&wlow` via queue `SP1IH` or `SP2IH` (depending upon the port specified by the user).

2. `S&wlow` transmits the message in a frame to the destination copy of `s&wlow`.

3a and 3b. `S&wlow` puts the message onto queue `APPLICATION` for processing by `s&wdisk` (step 3a), while simultaneously acknowledging the frame (step 3b).

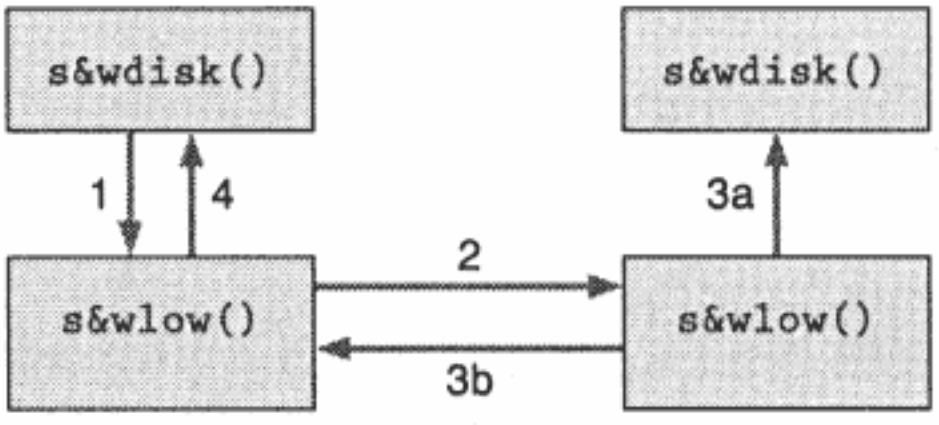

Figure 7.2
The message flow in the file transfer software.

4. S&wlow indicates to s&wdisk (via a completion code put onto queue APPLICATION) that the message has been received by the destination PC. At this point, s&wdisk can send another message.

7.4.3
Compiling the File Transfer Software

The file transfer example, s&wdisk.c, can be compiled and linked with s&wlow.obj, commkit.obj, srcalls.obj, and crc.obj, along with the various support routines, using the make utility:

```
C:\> make s&wdisk.exe
```

7.4.4
Running the File Transfer Software

Once created, s&wdisk. exe can transfer files between PCs connected by a null-modem RS-232-C cable; either serial port can connect to the other (i.e. 1 (or 2) to 1 (or 2)).

Since the stop-and-wait software is half-duplex, one PC must be designated as the source (of the file) and the other as the destination.

The source PC is invoked by typing s&wdisk followed by the line speed, the port (1 or 2), an *X* (signifying that this station is transmitting the file), and the name of the file to transfer, after the MS-DOS prompt. Error messages are issued if the line speed is incorrect, the wrong number of arguments is placed on the line, or the file does not exist.

For example, to transfer commkit.c to the destination PC on a 9600-bps line through port 2, type

```
C:\> s&wdisk 9600 2 X commkit.c
```

The transfer will continue until an error occurs or the file is completely transferred to the destination. Upon completion of the transfer, control returns to MS-DOS.

The destination station is initialized in a similar fashion; however, an *R* is used rather than an *X* to indicate that the station is receiving, and the file created and written to by s&wdisk is given the file name. Error messages are issued if the line speed is incorrect, the wrong number of arguments is placed on the line, or the file already exists.

For example, to create a file named commkit2.c with the information received from a 9600-bits-per-second line on port 1, type

```
C:\> s&wdisk 9600 1 R commkit2.c
```

Control is returned to MS-DOS when the receiving station receives the end-of-file signal from the transmitting station. If things go wrong, the receiving station can be aborted using CTRL-ALT-DEL.

7.5 Summary

This chapter examined file transfer and how it is achieved using a simple file transfer protocol running atop the Commkit stop-and-wait protocol. Other file transfer protocols are possible and will be discussed in subsequent chapters.

Protocols are used in almost all aspects of computing science. For example, consider the transfer of information between two DTEs involved in the file transfer. Rules are applied to the transfer of the message: the message must conform to a certain format; files cannot be accessed unless they are opened; and files can be accessed only according to the way they are opened.

7.6 Exercises

1. (B) Perform a file transfer with the line analyzer between the sending and receiving PCs. Identify the various fields in the data packets.

2. (B) Run `s&wex` as a receiver and `s&wdisk` in source modedoes the file transfer take place? Just because `s&wdisk` receives an acknowledgment for each record, does that mean that the file is being written to disk?

3. (B) Start `s&wdisk` in source mode on one DTE before starting `s&wdisk` in destination mode on a connected DTE. Does the transfer still work? Explain.

4. (B) Does the disk speed dictate when the destination `s&wdisk` software terminates?

5. (B) Try transferring an executable file. Does the transfer work? Explain. The line analyzer will help solve this exercise.

6. (C) Modify `s&wlow.c` so that executable files can be transferred; then transfer an executable file between the two PCs. (Before making any modifications, it is advisable to study the stop-and-wait protocol state diagrams in Figures 6.66 and 6.67.)

Does the transferred file function properly on the receiving PC? Since there is no end-to-end verification that the file transfer is successful, is this the only method of determining that the file was sent (and received) correctly?

7. (C) Modify `s&wdisk` so that the file name is supplied to the destination DTE by the source DTE. The software should still accept files with duplicate file names. *Hint*: add an option to the list of available options; see Table 7.2.

8. (C) Is it possible to transfer file attributes so that the newly created file looks "the same" as the original? The UNIX file functions `chmod()` and `stat()` offer some interesting possibilities. Now the tricky bit, can the file's original time of creation be kept?

9. (C) Modify `receive_and_write()` so that CTRL-C will abort the receive software.

Table 7.2 Hint for Exercise 7.

Command	Option	Data
Open file	O	Name of file
File data	D	One record from the file
Close file	C	Empty

10. (C) One limitation of the existing file transfer utility is its inability to transfer more than one file without returning to MS-DOS. Rewrite the file transfer software so that a series of files can be transferred without requiring the software to be restarted. For example:

```
C:\> s&wdisk 9600 1 X

file1.c
TRANSFER STARTED
TRANSFER COMPLETED

commkit.obj
TRANSFER STARTED
TRANSFER COMPLETED

s&wdisk.exe
TRANSFER STARTED
TRANSFER COMPLETED

CTRL C
C:\>
```

The file name can be obtained in one of two ways:

> The keyboard software in `low_level()` can assemble the file name until a carriage return is detected, at which point the string can be forwarded to `APPLICATION`. This involves changing both `s&wlow.c` and `s&wdisk.c`.
>
> The file name can be assembled in `read_and_send()`, one byte at a time, until a carriage return is detected. This involves only changing `s&wdisk.c`.

11. (C) Many line analyzers have the ability to record data displayed on the screen for subsequent playback. Since Commkit supports simultaneous file access and UART interrupts, try implementing a data capture routine on the line analyzer.

8
The Telephone System

8.1
Introduction

Consider the situation in which a small, local company grows in size and opens branch offices in a number of cities. For the company to function successfully, the people in the branch offices need to communicate with the head office, either by telephone, courier, or mail. In addition, the people in the branch offices may need to access information maintained on the computer(s) at the head office. The problem is, how can this be done? Surface mail and courier may be too slow for many applications. Ideally, the information will be transferred electronically, but given the distance limitations placed on RS-232-C (and most of the other standards discussed in Chapter 3), the required transfer cannot be directly supported.

What is needed is a communication facility that will allow the interconnection of two DTEs regardless of their locations (whether in the same building or on different sides of the world). The most obvious answer is, of course, the telephone system (or GSTN, General Switched Telephone Network), simply because it is the most readily available.

This chapter examines the telephone system and its use as a medium for the transmission of data, as well as the equipment needed to access it. The chapter also includes a detailed description of RS-232-C (ITU-T V.24) and the ITU-T recommendations for high-speed modems. The Commkit software demonstrates how the telephone system can transmit and receive data using a PC with either an external or internal modem.

8.2
The Telephone System

The telephone system was originally designed to transmit voice information from one subscriber, through a central, or end, office, to another subscriber.

The end office was (and still is) responsible for setting up, maintaining, and taking down the connection. The mechanisms for supporting these operations have changed substantially over the past century: from direct operator intervention, through a series of electromechanical devices (originally the Strowger step-by-step switch and eventually the cross-bar switch), to the digital switching of today. The future appears to be changing as rapidly as the past, with a move away from analog voice technology to worldwide digital transmission known as ISDN, or Integrated Services Digital Network.

The basic concepts associated with the telephone have not changed as drastically. Subscribers still take the telephone handset off-hook to establish a connection with the end office. The end office detects the off-hook condition (current is drawn) and responds with a dial tone. The number selected by the subscriber (either through dialing, which breaks the current at precise intervals, or by creating tones at precise frequencies) is recorded by the end office. Once the number is obtained, the end office attempts to make the connection with the other party. If the other party's line is not engaged (i. e., busy), the end office generates a signal that causes the telephone to ring. When the call is answered, a circuit is then dedicated for the duration of the call. The circuit is freed by either party putting the handset on-hook.

The end office is designed to handle local calls for a limited number of subscribers (in North America, roughly 10,000 per end office). Since most telephone companies have more than 10,000 subscribers, the telephone system consists of thousands of end offices. Although every subscriber's telephone is connected to an end office through a two-wire local loop, the end office has sufficient equipment to handle only about 10 percent of attempted calls (up to 20 percent in some locations where there is a great deal of business traffic).

Not all calls are local to the subscriber's end office; some are intended for subscribers connected to other end offices. These are long-distance calls and require special signaling to allow one end office to determine that the call is for another one. The call is routed from the subscriber's end office through one or more circuits to the end office of the specified party. In North America this is achieved by prefixing the dialed party's number with a 1 and possibly an area code.

The technology used to create and maintain the circuit is known as circuit switching. The paths taken through the telephone system can vary depending upon a number of factors, such as time of day and day of week. Once the circuit is established, it remains in place until the call is freed. Figure 8.1 illustrates the basic structure of the present North American telephone system, consisting of some 150 million subscribers.

The circuit eventually established for the call is dictated primarily by the number of potential circuits between the two end offices and the volume of traffic on the telephone system at the time the call is placed. Ideally, the

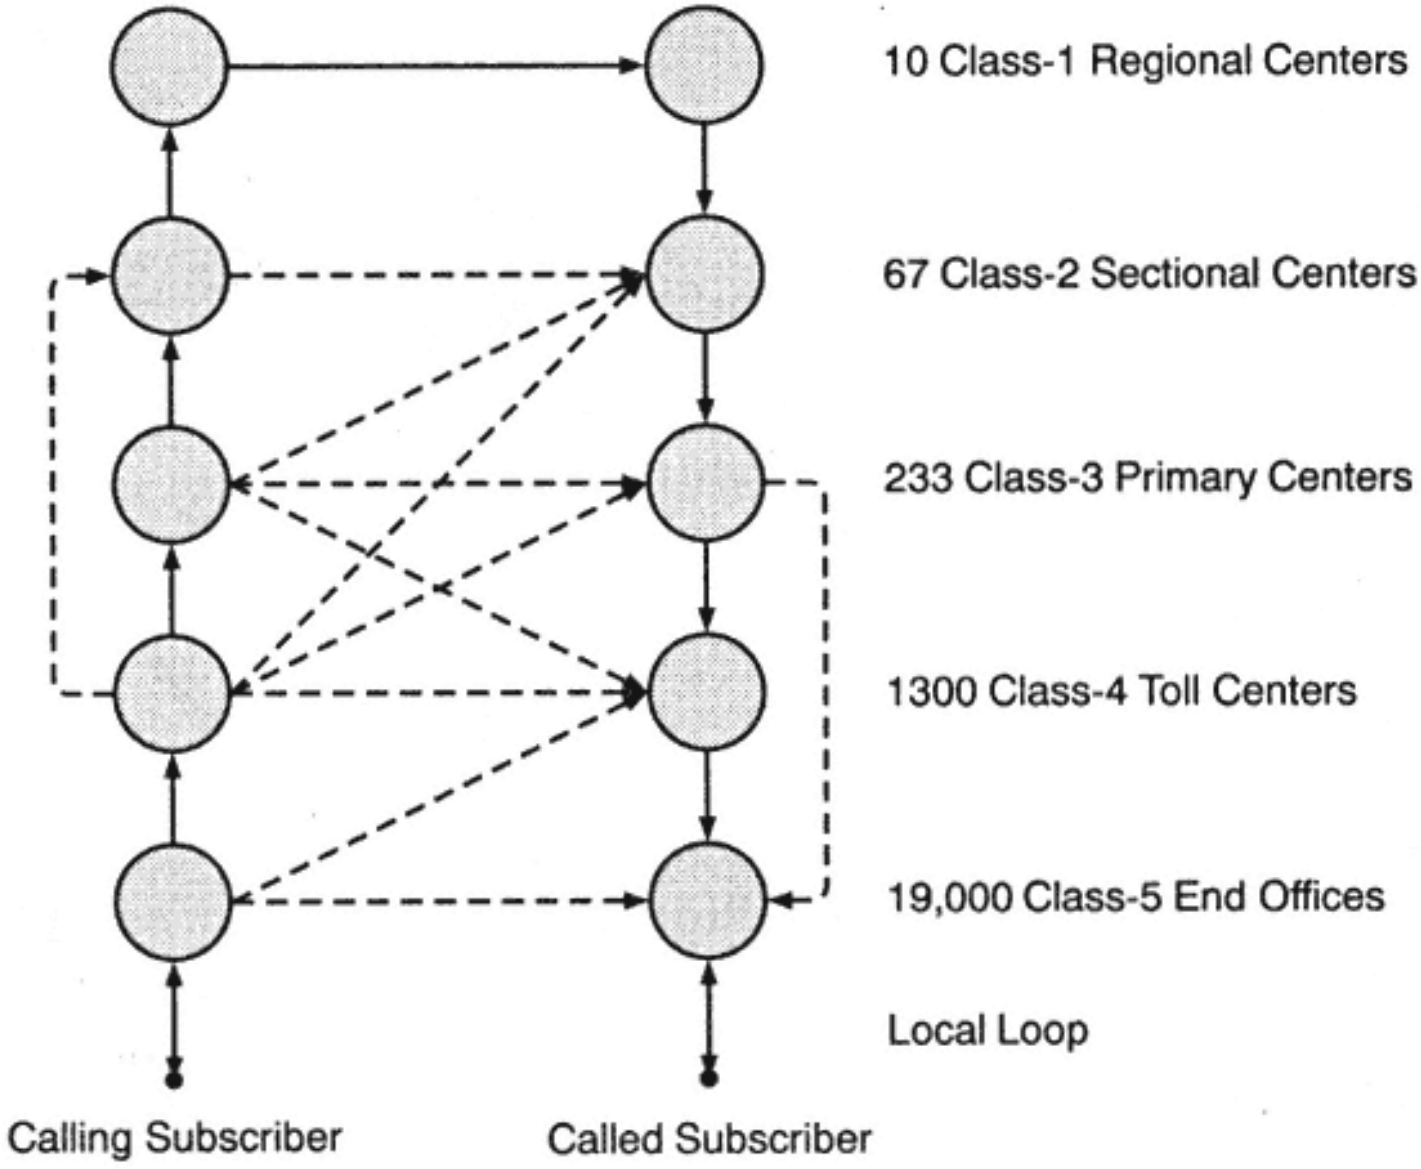

Figure 8.1
Basic structure of North American telephone system.

call will be routed across the shortest path connecting the two end offices; but sometimes that path may be fully utilized by other calls. Recognizing this, telephone engineers have developed a network of trunks that allow the subscriber calling to reach the subscriber called through one of a number of different paths. Some heavily used paths, represented by dashed lines in Figure 8.1, are known as high-usage trunks and can be thought of as shortcuts across the network between end offices or centers. If the path across a high-usage trunk is unavailable, traffic overflows are placed on the next available high-usage trunk or, as a last resort, on a final trunk (represented by a solid line in the Figure 8.1). (The present hierarchical structure is evolving to support a dynamic routing structure not restricted to preplanned high-usage trunks and final trunks.)

The equipment connecting the subscriber's telephone to the end office is typically analog and designed for voice frequencies in the 300- to 3400-Hz range. As was shown in previous chapters, signals sent across wires are subject to attenuation and can eventually die out entirely. To avoid this problem, signals are "boosted" using some form of signal regeneration. Prior to the development of digital signaling, the signals on a telephone circuit were boosted using amplifiers. If the signal was subject to noise, both the voice and the noise were amplified; in the worst case, all that was heard by either subscriber was the noise!

Digital technology has helped eliminate many of these problems. In dig-

Table 8.1 The North American telephone multiplexing hierarchy.

Designation	Data Rate	Circuits
T1	1.544 Mb/s	24 × 64-kb/s digital voice
T2	6.312 Mb/s	4 × 1.544 Mb/s T1
T3	44.736 Mb/s	7 × 6.312 Mb/s T2
T4	274.176 Mb/s	6 × 44.736 Mb/s T3

ital end offices, the analog voice signals are converted to digital ones using an analog-to-digital converter. The digital voice signals are converted back to analog (using a digital-to-analog converter) at the receiving subscriber's end office. Digital signals sent between center offices are multiplexed into groups of signals using a technique known as time division multiplexing, or TDM. Time division multiplexing takes a number of different telephone circuits, samples them at precise intervals, and sends the value of the sample over a communication channel to another end office or center. The receiving office or center samples the communication channel at the same rate as the transmitter and routes the resulting sample to either the subscriber or another office or center. Nyquist's sampling theory states that sampling at *twice* the highest signal frequency yields a sample that contains all the information of the original signal. Since voice data falls in the range of 0 to 4000 Hz, sampling the channel 8000 times a second captures all of the voice data.

In North America, a multiplexing hierarchy has been developed to support circuits between end offices. Table 8.1 illustrates the North American hierarchy, while similar ones exist for telephone transmission rates in Europe and Japan.

8.3 Connecting to the Telephone System

As demonstrated at the start of this chapter, the obvious solution to communicating over distances that cannot be supported by RS-232-C and other standards is to use the telephone system. However, several problems must be addressed before the communication can take place, notably:

The telephone operates at frequencies in the range of 300 to 3400 Hz, while a DTE generates voltages of extremely low frequencies that a telephone will not detect.

How can the connectors found on the back of a DTE be connected to the telephone system?

In short, it is not possible to hold a telephone handset up to the back of a DTE in the hope that the DTE's signals will be detected and transmitted. A device is needed to convert (or modulate) the DTE's electrical signals into a frequency that can be used by the telephone. This device must also be able

to convert (or demodulate) the telephone frequency signals into electrical impulses that can be understood by the DTE.

The device that converts these signals into telephone frequencies and vice versa is known as Data Communications Equipment or Data CircuitTerminating Equipment (abbreviated DCE). More commonly, the DCE is referred to as a modem, indicating the functions the DCE performs: *mo*dulating the voltage into a frequency and *dem*odulating the frequency into a voltage.

8.3.1
Modems

Whenever a telephone call is made, two stations are involved: the one that originates the call and the one that answers it. The same terminology is applied to pairs of modems involved in a communication: the one where the call originatesthe originating modem; and the one where the call is answeredthe answering modem.

Modems can be purchased as originate-only (that is, they can only place calls), answer-only (they can only receive calls), or originate/answer (they can either originate or answer). An originating modem can communicate only with an answering modem; that is, originateoriginate and answeranswer combinations are not possible.

8.3.2
Modem Signaling

The telephone works roughly in the frequency range of 300 to 3400 Hz. The modem, using the telephone system, therefore has about 3000 Hz in which to modulate and demodulate the binary information. Three basic techniques can be used to encode binary information on the telephone system, all of which are now considered.

Amplitude Modulation

In amplitude modulation (or AM), a mark (1) is represented by one amplitude of a carrier frequency and a space (0) is represented by another (see Figure 8.2).

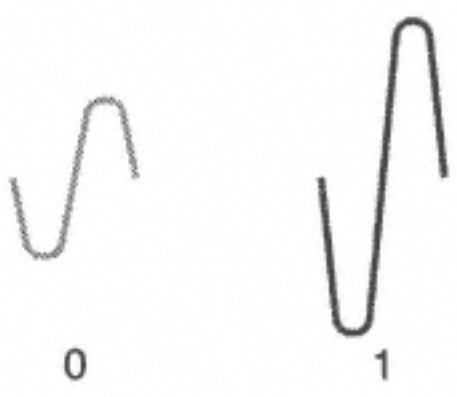

Figure 8.2
Amplitude modulation.

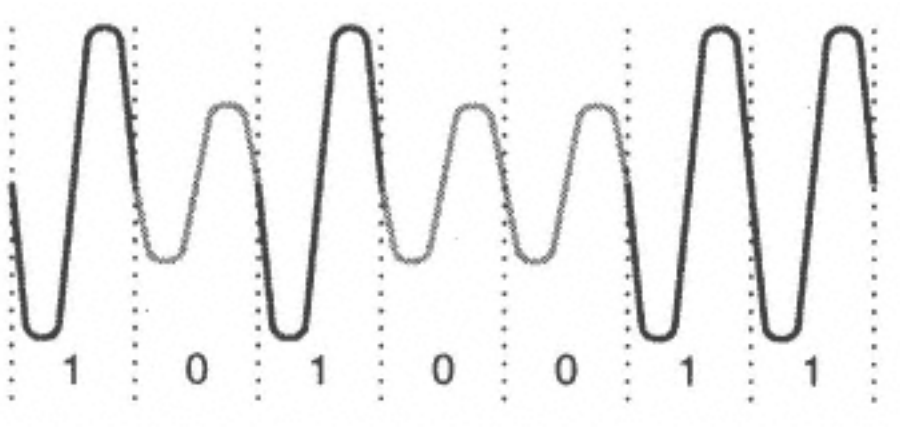

Figure 8.3
Amplitude modulation of the byte *S*.

The binary information to be transmitted is keyed (or switched) between these two amplitudes by the modem. For example, Figure 8.3 shows how the ASCII byte *S* (with bit pattern `1010011`) can be transmitted. The pattern represents the byte *S* being sent as "pure" data. Had the byte been sent asynchronously, the start, parity, and stop bits would also have been included in the signal.

Since the amplitude modulation signal is prone to signal attenuation, modems using amplitude modulation are not in common usage other than at very low line speeds, such as 300 bits per second.

Frequency Modulation

Frequency modulation (or FM) differs from amplitude modulation in that the amplitude of the carrier is fixed but its frequency changes in accordance with the binary data being transmitted. In other words, a mark (1) is sent at one frequency and a space (0) at another (see Figure 8.4).

Frequency modulation is also known as frequency-shift keying, or FSK. Figure 8.5 shows a possible frequency modulation of the byte *S*. FSK modems are generally designed to operate at line speeds in the range of 300 to 1200 bits per second.

Phase Modulation

In phase modulation (or PM), the carrier's frequency and amplitude are kept constant but its *phase* is shifted to represent a mark (1) or a space (0). For example, in Figure 8.6 mark and space are shifted 180° to one another.

When binary signals are encoded this way, it is known as phase coher-

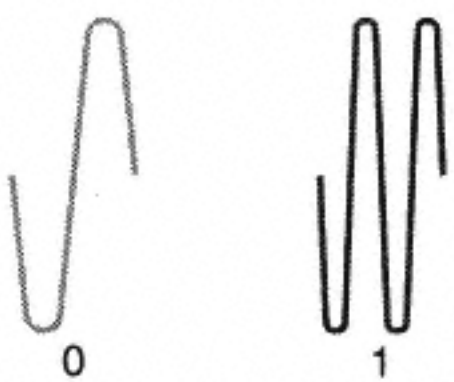

Figure 8.4
Frequency modulation.

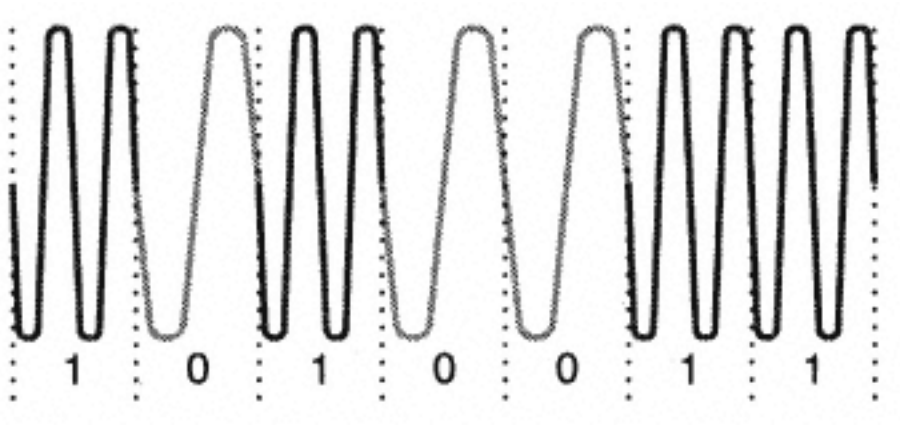

Figure 8.5
Frequency modulation of the byte S.

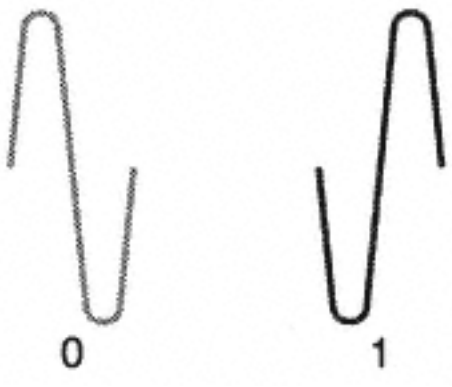

Figure 8.6
Phase modulation.

ent phase-shift keying, or phase coherent PSK. In Figure 8.7, the byte *S* is transmitted using PSK.

In phase coherent PSK, the modem maintains internal reference signals that are compared against the incoming signal. Phase coherent PSK modems suffer from several drawbacks, including their need for complex demodulation circuitry and their susceptibility to random phase changes in the signal.

A variation on phase coherent PSK is differential PSK (or DPSK). DPSK differs from phase coherent PSK in that the signal used to represent a bit is *relative* to the last bit signal. For example, a phase shift of 90° could indicate a space, while a phase shift of 270° could indicate a mark. Figure 8.8 shows the byte *S* encoded with DPSK using the above shifts.

A DPSK modem need only determine the magnitude of the phase shift to obtain the value of the bit.

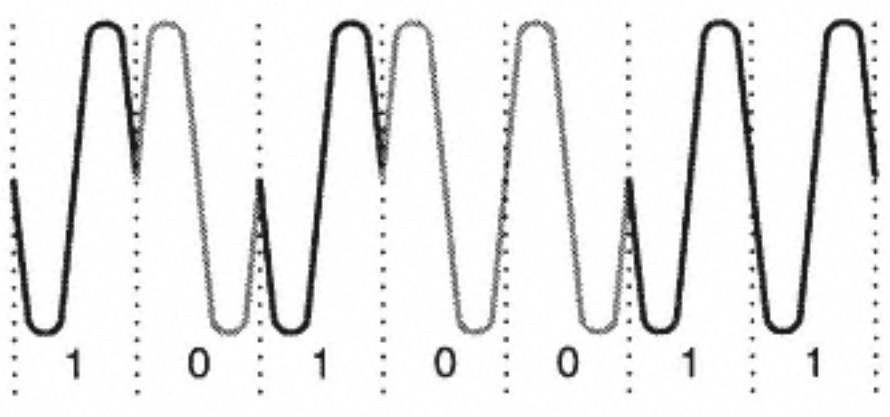

Figure 8.7
Phase modulation of the byte *S*.

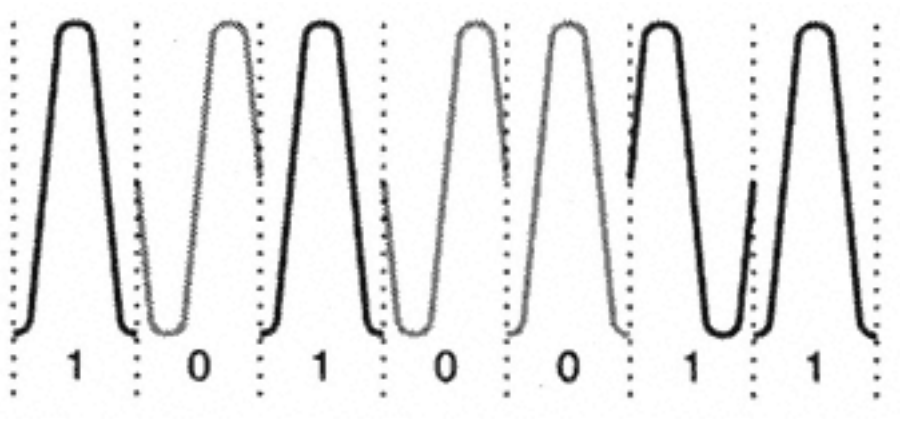

Figure 8.8
Differential PSK encoding of the byte *S*.

Example

The modem offers a limited bandwidth of about 3000 Hz; if the entire frequency is devoted to one DTE or the other, the result is a simplex communication (or half-duplex if a protocol is in place to allow the channel to be turned around).

Achieving a full-duplex communication requires that the available frequency (3000 Hz) be divided in two: one half for transmission and the other half for reception. This raises a problem: there are two modems, and if they both transmit at the same frequency, their signals will collide and communications will cease. Some type of agreement is necessary to ensure that one modem will use one set of frequencies and the other modem another set of frequencies.

Fortunately, it is possible to assign frequencies based upon whether the modem is originating or answering the call (hence the reason for allowing only originating modems to communicate with answering modems). The originating modem will transmit on frequency *X* and receive on frequency *Y*, while the answering modem will transmit on frequency *Y* and receive on frequency *X*.

As an example, Figure 8.9 shows the modem frequency assignment for a 300-bps Bell-103 modem; the frequencies depend upon whether the modem is in originate or answer mode.

The ITU-T equivalent 200-bps modem (ITU-T Recommendation V.21) uses the frequencies shown in Figure 8.10.

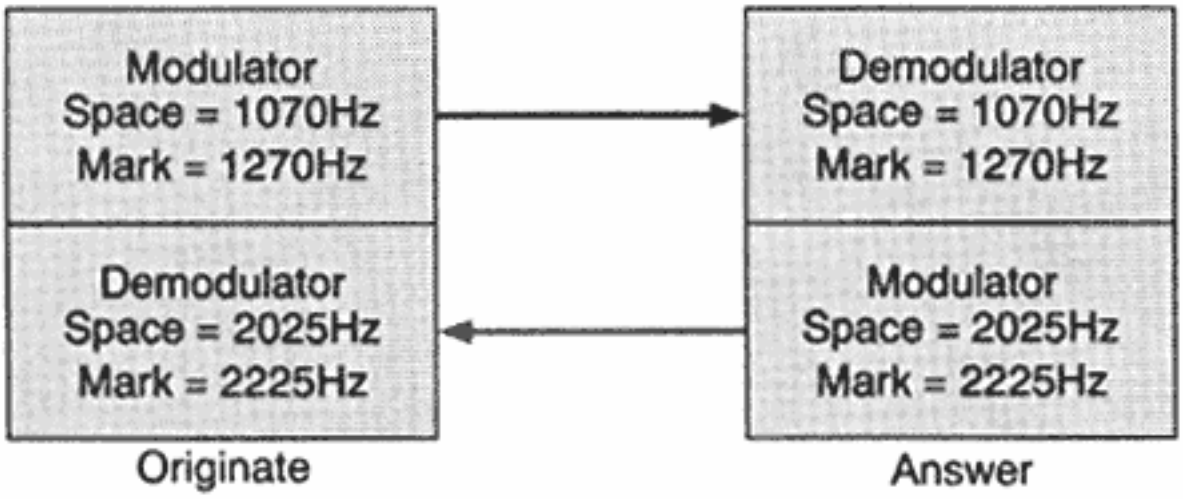

Figure 8.9
Frequency assignment for a 300-bps Bell-103 modem.

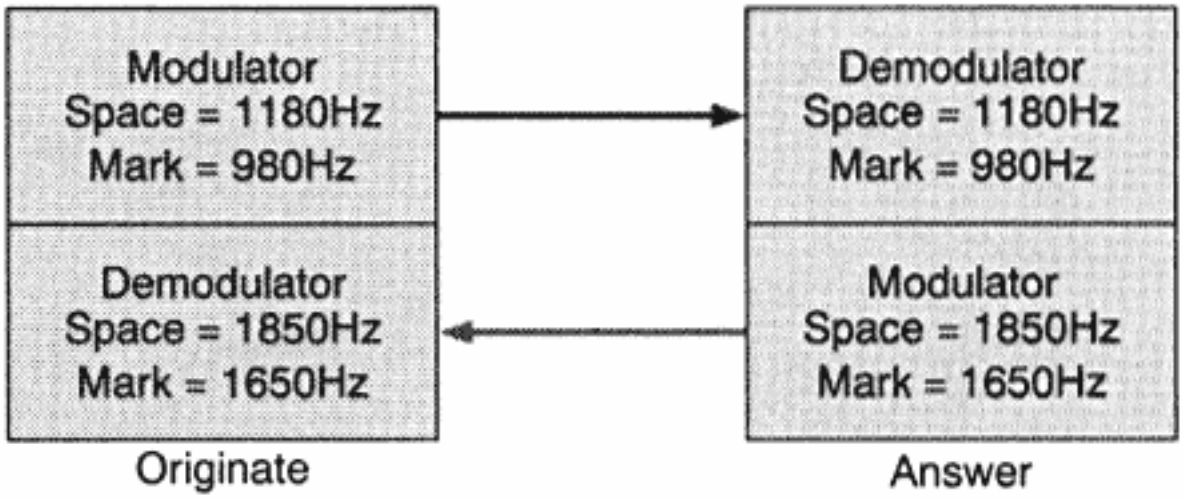

Figure 8.10
Frequency assignment for a ITU-T 200-bps modem.

Bits-per-second vs. Baud

Up to this point in the text, all references to line speed have been in terms of bits-per-second. However, in reading through the literature one finds that modems are often described in terms of baud rather than bits per second. Baud (named in honor of Emile Baudot, inventor of the Baudot code) is defined as a unit of signaling speed equal to the number of discrete conditions or signal events per second.1

Baud is equivalent to bits per second if each signal event is equal to exactly one bit. As an example, 300 bps is the same as 300 baud if there is one signal event every one-three-hundredth of a second. If the channel uses 2400 Hz (cycles per second) to indicate a mark, the number of cycles required to represent a mark is

$$\frac{1}{300}\ \text{sec} \times 2400\ \text{Hz} = 8\ \text{cycles}$$

A bit encoded at 2400 Hz is shown in Figure 8.11.

Similarly, 600 bps is equivalent to 600 baud at 2400 Hz if one bit is sent every 4 cycles. If the modem can be designed to recognize one bit per cycle, it can support 2400 bps. Since the maximum frequency the telephone system allows is about 3400 Hz, the maximum line speed that can be achieved if one bit is equivalent to one baud is 3400 bps. However, the U(S)ART can be programmed to much higher speeds than 3400 bps and still use a modem: the question is, how?

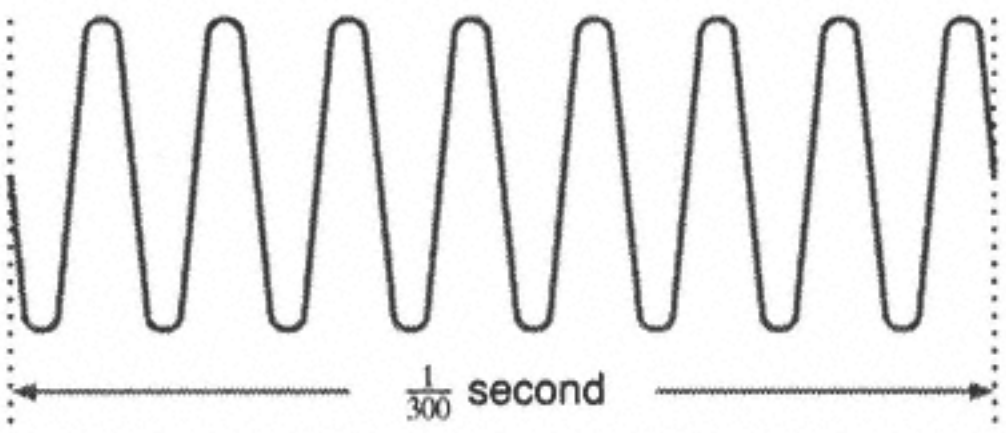

Figure 8.11
A 300-bps bit at 2400 Hz.

1J.E. McNamara. *Technical Aspects of Data Communications*. 2nd ed. Digital Press, 1982.

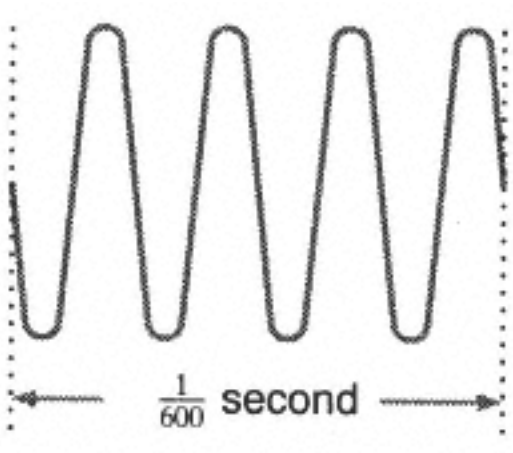

Figure 8.12
One signal event (baud) at
2400 Hz and 600 baud.

One answer is to use phase shifting. For example, at 600 baud there are 600 "signal events" occurring each second. Figure 8.12 shows that at 2400 baud there are 4 cycles per baud, and Figure 8.13 shows that by shifting the phase of the signal by 90°, four distinct patterns emerge, allowing the assignment of four different bit combinations.

Now each signal event can represent *two* bits; two bits per baud is known as a dibit. The phase shift allows the 600-baud channel to support 1200 bps. For example, transmitting the 7-bit ASCII character *S* on a 600-baud channel at 2400 Hz produces the signal shown in Figure 8.14 (assuming phase coherent PSK). Note that an even number of bits are required and, in this case, the eighth bit is sent as a 0. In an asynchronous transmission, the entire frame is sent (i.e., start, stop, parity, and data bits). If the resulting number of bits is odd, the next start bit or an idle mark will be encoded into the last dibit. Similarly, in a synchronous communication a dibit may contain the last bit of one data byte and the first bit of the next data byte.

The Bell 212 is a full-duplex modem that supports either 300-bps or 1200-bps communications. In 300-bps mode, the modem generates FSK signals and is compatible with the Bell-103 modem described above. 1200 bps is achieved

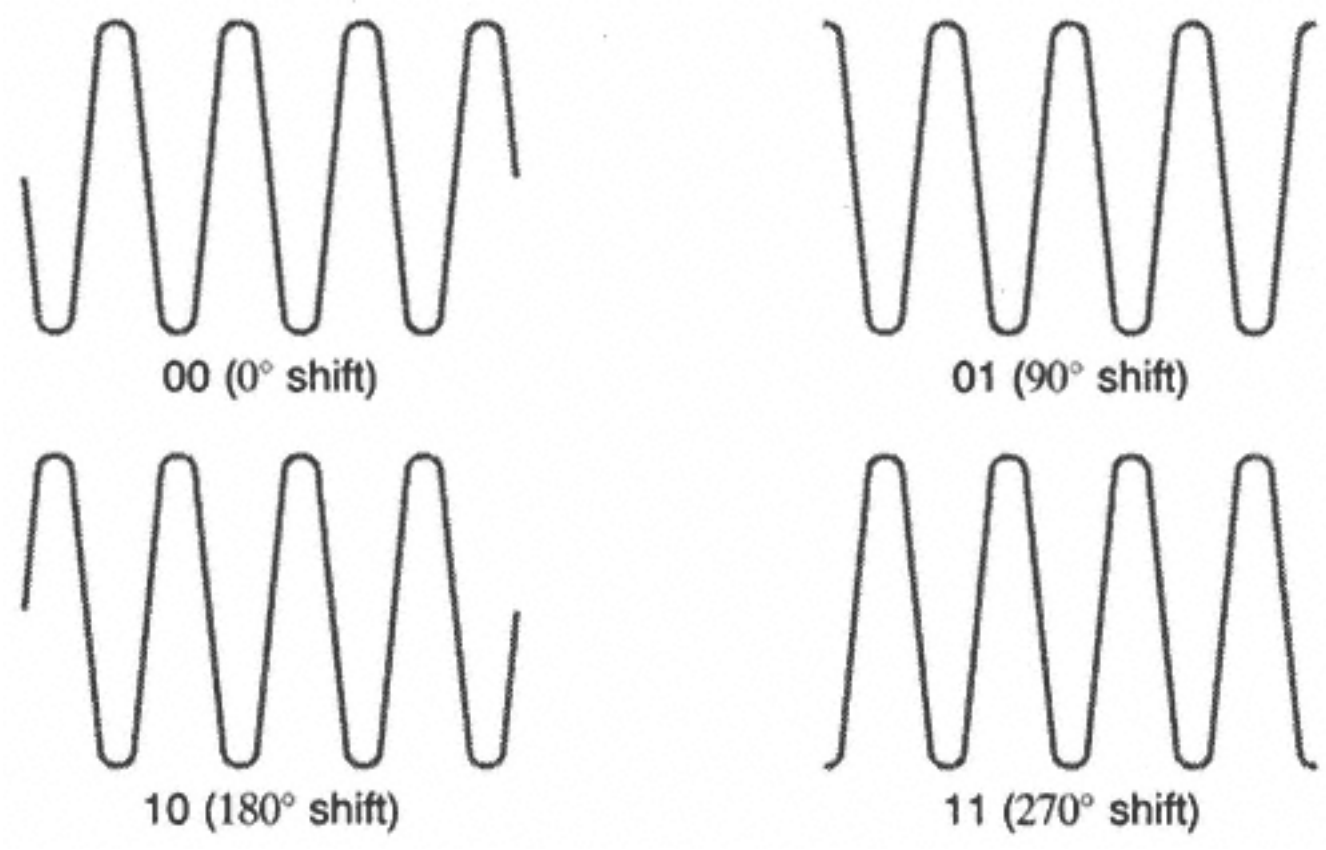

Figure 8.13
A 90° phase shift allows two bits per baud.

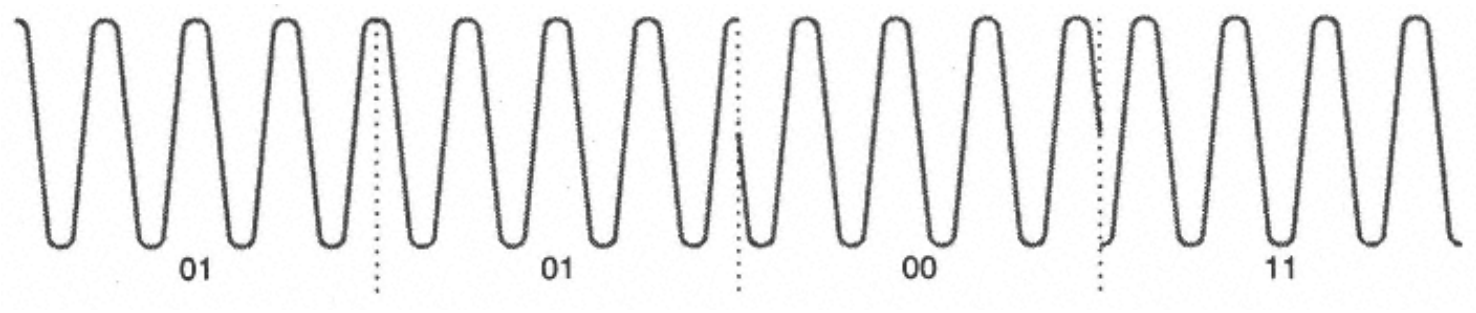

Figure 8.14
The byte *S* encoded at 2400 Hz and 600 baud with a 90° phase shift.

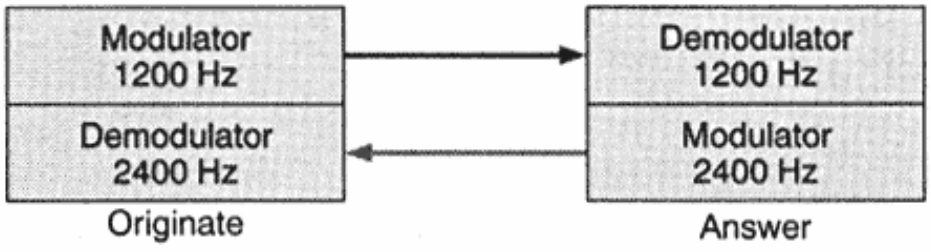

Figure 8.15
1200-bps Bell-212 frequency assignments.

using PSK operating at 600 baud; as shown in Figure 8.15, there is a single transmission frequency and a single reception frequency.

The Bell 212 achieves 1200 bps on a 600-baud channel by encoding the data into dibits using differential PSK; dibit values are shown in Table 8.2.

8.4 RS-232-C (ITU-T V.24)

RS-232-C (ITU-T V.24) was introduced in Chapter 4 as a mechanism to allow pairs of DTEs to communicate. However, the original purpose of RS-232-C was to allow DTEs to connect to DCEs, thereby permitting communications across telephone systems. This section examines how RS-232-C supports DTEDCE connections.

The signals between the DTE and the DCE are digital 0s and 1s, while the signals on the telephone network are, depending upon the modem, the analog signals that were discussed in the previous section. Figure 8.16 depicts the relationship between RS-232, the DTE, and the DCE.

Table 8.2 Bell-212 dibit values.

Dibit	Phase Differential
00	135
01	45
10	+45
11	+135

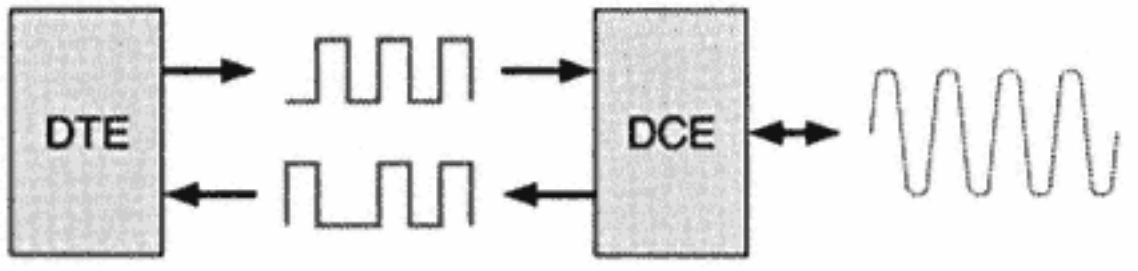

Figure 8.16
A DCE supports'digital and analog signals.

RS-232-C requires a minimum of three connections between the DTE and the DCE: signal ground, transmission, and reception. To ensure that all manufacturer's equipment (DTE or DCE) can communicate, RS-232-C defines both its signal levels and the physical wiring between the DTE and the DCE. RS-232-C connectors support either 25 or 9 pins; each pin performs a specific task and is given a unique label. All signals are defined in relation to the DTE; that is, signals are sent *to* or received *from* the DCE.

8.4.1
Required Connections

Although RS-232-C supports a 25-pin connector, only three connections are required for the transmission and reception of data (two circuit labels are given: the first for RS-232-C, the second for the ITU-T V.24 100-Series interchange circuit definitions):

Pin 2. Transmitted Data (to DCE). Circuit BA (ITU-T: 103). Data generated by the DTE is sent to the DCE on this circuit. The circuit is kept in the mark state at all times when there is no transmission in progress.

Pin 3. Received Data (from DCE). Circuit BB (ITU-T: 104). Signals are generated on this circuit by the DCE based upon the signals received from the remote DCE. The circuit is kept in the mark state at all times when there is no transmission in progress.

Pin 7. Signal ground or common return (there is no direction on this circuit). Circuit AB (ITU-T: 102). This circuit establishes the common ground between the DTE and DCE. It is used as the reference for determining whether a signal is a mark or a space.

8.4.2
Minimum Line-Control Requirements

Transmit data, receive data, and signal ground are the minimum number of connections required between the DTE and DCE to ensure data transfer. The following circuits are used to control access to the telephone channel:

Pin 22. Ring Indicator (from DCE). Circuit CE (ITU-T: 125, "Calling indicator"). Set by the DCE when an incoming ringing signal has been detected. The signal is *on* during a ring and *off* between rings and at all other times when ringing is not being received.

Pin 20. Data Terminal Ready, more commonly DTR (to DCE). Circuit CD (ITU-T: 108/2). Used by the DTE to control the DCE's access to the communication channel. The *on* condition causes the DCE to connect to the telephone channel (this can be set at any time, either before Ring Indicator is set or during the ringing sequence).

The *off* condition causes the DCE to disconnect from the communication channel. It is important to note that because most U(S)ARTs are double buffered, one or more bytes may still be in the UART when the *off* condition is to be raised. To avoid losing these bytes, the software should be written so that it waits for an indication that the last byte has been sent (on the PC's 8250 UART, the signal is Transmit Shift Register Empty) before setting DTR *off*.

Pin 8. Received Line Signal Detector, more commonly carrier detect (from DCE). Circuit CF (ITU-T: 109, Data Channel Received Line Signal Detector). Indicates that the local DCE has detected a valid line signal from the remote DCE. The circuit is *on* when the signal meets the signal levels specified by the type of DCE; it is *off* when the signal no longer meets the required levels.

Pin 6. Data Set Ready, more commonly DSR (from DCE). Circuit CC (ITU-T: 107). An indication as to the status of the local DCE. *On* indicates the following conditions:

1. The local DCE is connected to a communication channel.

2. The local DCE is not in test, talk, or dial mode.

3. Has reached the stage where a call can be established (see Section 8.4.3).

An *off* indicates that the DTE is to ignore the signals from the DCE.

8.4.3
Full-Duplex Communications

In a typical full-duplex communication, each DTE is given half of the carrier, thereby allowing simultaneous transfer of information. Before a communication can take place across the telephone network, the call must be set up (or established) by the modems (one originating the call, the other answering it). Although the exact sequence can vary from manufacturer to manufacturer, the steps required to set up, maintain, and eventually clear a full-duplex communication are essentially as follows.

Originating Modem

At the originating modem, the call is placed by a user following these steps or a variation on them:

1. The user dials the digits of the answering modem.

2. The telephone at the answering modem rings. The answering modem responds with a carrier once the telephone has been answered (there can be a delay in this step if the answering modem requires operator intervention; see below).

3. The user presses the Data/Talk button.

4. The telephone is hung up (this does not clear the call, since the modem is in control of the telephone line).

Within the DTE itself, the following algorithm can be applied:

1. Wait for carrier and Data Set Ready (DSR).

2. Once detected, respond with Data Terminal Ready (DTR).

3. Commence and maintain the data transfer.

Note that most modems have been designed to allow steps 1 and 2 to be reversed.

Answering Modem

At the answering modem, call setup can be achieved by the DCE and DTE or by operator intervention. If the call is answered by the operator, the steps are as follows:

1. The telephone rings and is answered by the operator.

2. The DCE is in answer mode (to allow the transfer of data).

3. The operator presses the Data/Talk button.

4. The carrier is generated.

5. The operator hangs up the handset.

Internally, the software must wait for Carrier Detect and Data Set Ready. The subsequent steps are the same as those of the originating modem.

It is also possible to have the DTE answer the call in conjunction with the DCE. In such situations, the following steps occur:

1. The DCE detects a ring and signals the DTE with a Ring Indicator.

2. The DTE responds by asserting Data Terminal Ready.

3. The DCE generates a carrier and signals Carrier Detect and Data Set Ready to DTE.

4. The data transfer is performed as usual.

Call Clearing

The call can be cleared by either end through one of a number of methods:

The Data/Talk button on the modem can be pressed, resulting in the carrier being dropped. This causes Carrier Detect and Data Set Ready to drop, allowing the software to determine that the call has been cleared.

The software can stop asserting Data Terminal Ready. When the modem detects that DTE is no longer active, the carrier is dropped.

The carrier (from the remote DCE) drops. When the modem detects this (usually a signal that the remote has disconnected), Carrier Detect and Data Set Ready are dropped, once again allowing the software to clear the call.

In all of the above cases, the modem is usually designed to wait a period of time (on the order of 50 milliseconds) to ensure that the event has actually occurred. For example, a modem detecting a loss of carrier waits

50 milliseconds before clearing the call; should the carrier return during this period, the call is not cleared.

8.4.4
Half-Duplex Communications

A half-duplex communication requires that a single modem generate a carrier at any one time and differs from a full-duplex communication in that once one DTE has finished transmitting, the line can be "turned around," so that the other DTE can transmit (see Figure 8.17).

This raises the issue of how a modem "knows" when to turn the line around, start (or stop) transmitting, and stop (or start) receiving. Since RS-232-C is not an end-to-end protocol (that is, no information other than DTE data is passed between the two modems), the signal to perform line turnaround must come from another source.

The solution is to use a higher level of protocol to which both DTEs agree. For example, the half-duplex Bisync protocol defines which DTE can be transmitting at any one moment. As shown in Figure 8.18, certain agreed-upon control sequences cause the DTEs to exchange control of the line.

Since the modem is responsible for generating the carrier, it must be signaled when to start (and stop) carrier generation. RS-232-C defines two other signals that allow the DTE and modem to achieve line turnaround:

Pin 4. Request to Send, commonly known as RTS (to DCE). Circuit CA (ITU-T: 105). This signals the local DCE that the local DTE has data to send. However, the DTE cannot start transmitting at this point, but must wait for the DCE to indicate clear to send.

Pin 5. Clear to Send, commonly known as CTS or "Ready for Sending" (from DCE). Circuit CB (ITU-T: 106). Indicates that the local DCE is able to send data to the remote DCE. *On* indicates to the local DTE that data can be sent; *off* indicates that the local DCE is not prepared to transmit. Once the local DTE detects Clear to Send, it can start to transmit.

In a full-duplex communication, Clear to Send can be wired to the carrier

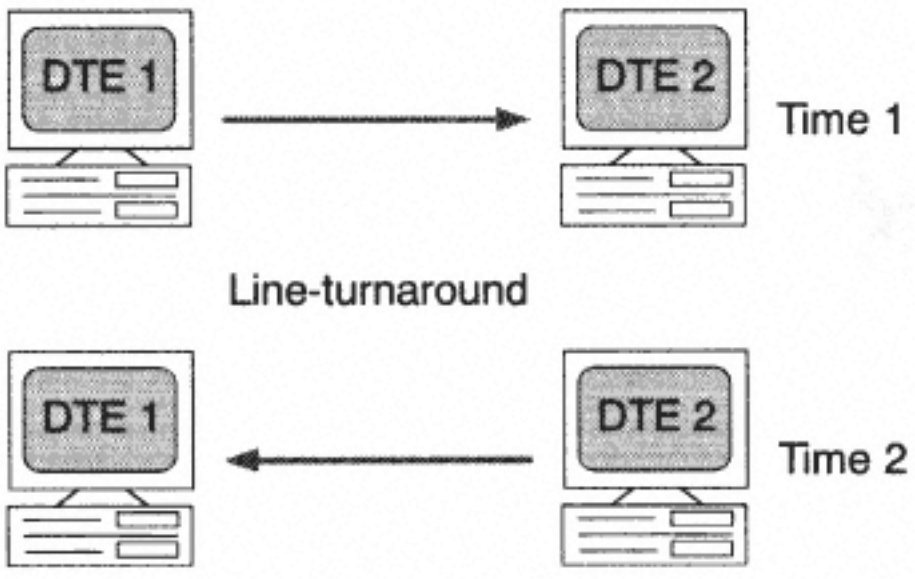

Figure 8.17
Half-duplex communications require line turnaround.

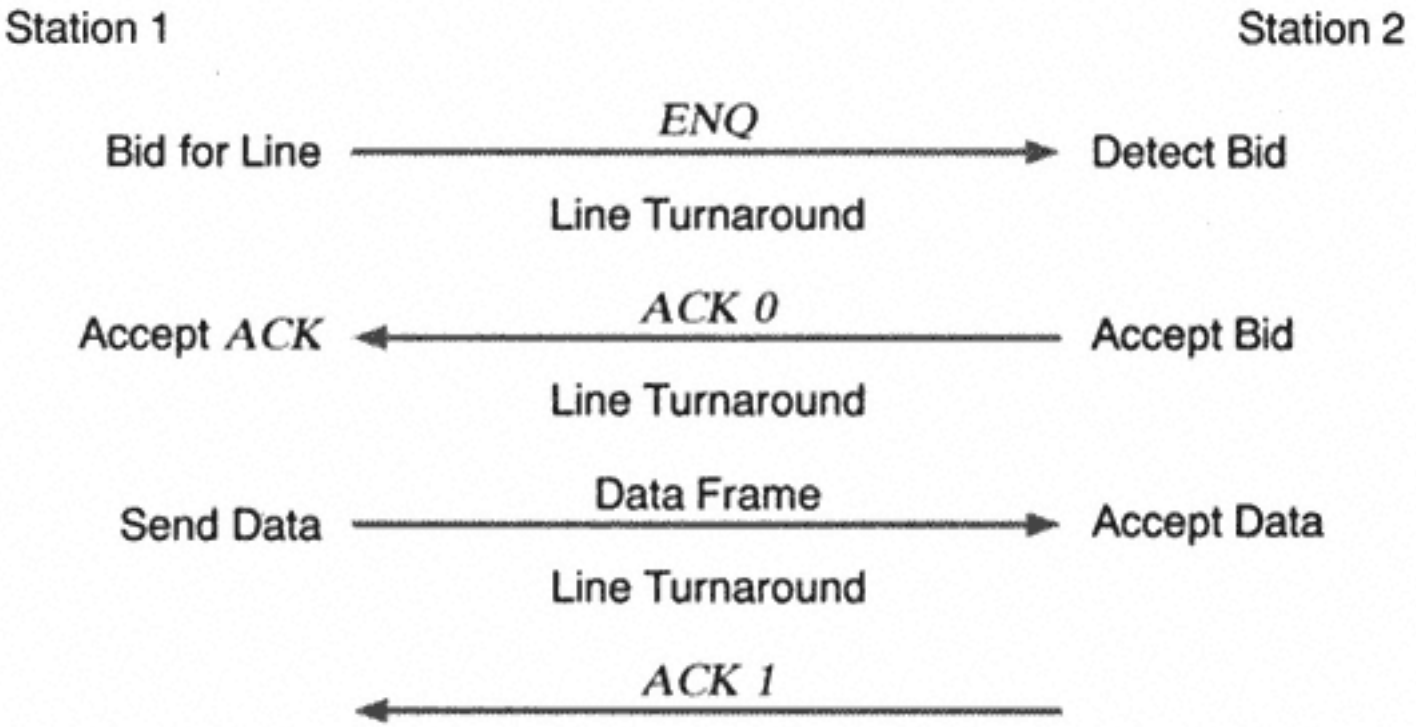

Figure 8.18
Using a higher-level protocol to control line turnaround.

detect signal, meaning that as long as there is a carrier, it is clear to send data (this is known as CB-CF common) and is shown in Figure 8.19. In some full-duplex modems, the Request to Send signal from the DTE is ignored, since it is assumed that the DTE will always be in a state that allows transmission.

In a typical half-duplex RS-232-C application, the DTE sends a Request to Send to the DCE, to which the DCE responds with Clear to Send. However, since RS-232-C is not an end-to-end protocol, the local DCE has no indication of whether the remote DCEDTE is ready to accept data (remember, the Clear to Send signal is simply between the local DTE and the local DCE). To avoid having the local DTE transmit data before the remote DCEDTE is ready to accept it, the Request to Send can invoke a timer that, upon expiring, causes Clear to Send (this is known as CB-CF separation) and is illustrated in Figure 8.20.

When the line is to be turned around, each DTE is responsible for certain actions to ensure that the communication can continue:

DTE about to receive. The DTE that is about to receive (i.e., the former transmitter) must signal its modem to stop generating the carrier signal; this is achieved by dropping Request to Send. This DTE then waits for its modem to detect the carrier from the remote DCE, at which point carrier detect is signaled. Data can then follow from the remote DCE.

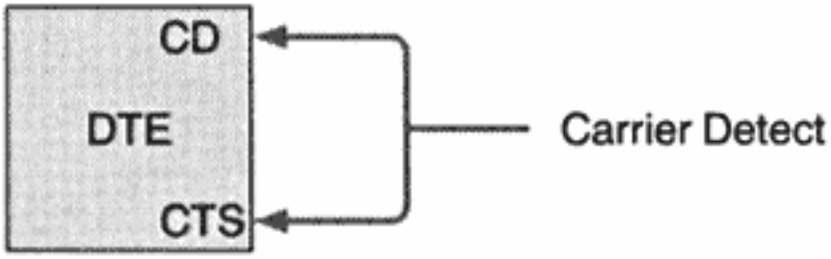

Figure 8.19
CB-CF common.

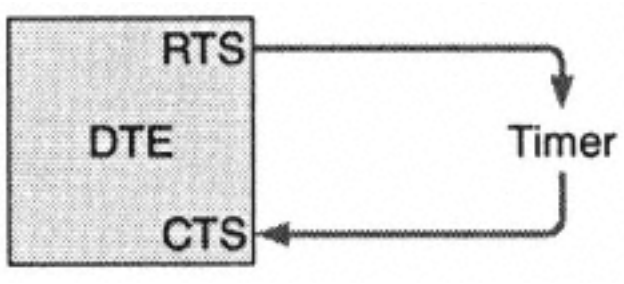

Figure 8.20
CB-CF separation.

DTE about to transmit. The DTE about to transmit (i.e., the former receiver) signals its modem to start generating the carrier signal by raising Request to Send. After a given time (see above), the modem signals Clear to Send to the DTE, which can then start to transmit.

8.4.5
Miscellaneous Connections

In addition to the nine circuits already described, there are a number of other RS-232-C circuits used to support a variety of communication situations. In general, these circuits are employed only in exceptional circumstances.

Secondary Communication Channel

RS-232-C supports a "Secondary Communication Channel" or reverse channel (ITU-T: Backward Channel) that is intended for the transmission of supervisory or error-control signals, typically in a half-duplex communication. The reverse channel runs in the direction opposite to the main transmission, thereby allowing the receiving DTE to signal without turning the line around. The circuits associated with the Secondary Communication Channel are

Pin 14. Secondary Transmitted Data (to DCE). Circuit SBA (ITU-T: 118). Equivalent to circuit BA, except that SBA is used by the secondary channel.

Pin 16. Secondary Received Data (from DCE). Circuit SBB (ITU-T: 119). Equivalent to circuit BB, except that SBB is used by the secondary channel.

Pin 12. Secondary Received Line Signal Indicator (from DCE). Circuit SCF (ITU-T: 122). Equivalent to circuit CF (carrier detect), except that SCF is used by the secondary channel.

Pin 13. Secondary Clear to Send (from DCE). Circuit SCB (ITU-T: 121). Equivalent to circuit CB (Clear to Send), except that SCB is used by the secondary channel.

Pin 19. Secondary Request to Send (to DCE). Circuit SCA (ITU-T: 120). Equivalent to circuit CA (Request to Send), except that SCA is used by the secondary channel.

Other Signals

Pin 1. Protective Ground. Circuit AA (No equivalent ITU-T circuit). This circuit is grounded to the DTE.

Pins 9, 10. Reserved for testing purposes.

Pin 23. Data Signal Rate Selector. This signal permits the changing of the data signal rate: if *on*. the higher speed supported by the DCE is chosen; if *off*. the lower speed is chosen.

This signal can be set by either the DCE or the DTE. If the signal direction is to the DCE, pin 23 is circuit CH (ITU-T: 111); otherwise, it is circuit CI (ITU-T: 112), from the DCE.

Pin 21. Signal Quality Detector (from DCE). Circuit CG (ITU-T: 110). If the telephone channel appears to be error free, this signal is set *on*. However, if the DCE detects an error on the telephone channel, it is turned *off*.

Pin 17. Receiver Signal Element Timing (from DCE). Circuit DD (ITU-T: 115). If the DTE does not support a receive clock that permits it to determine the center of each bit received, the clocking signal must be supplied by the DCE. This circuit allows the local (receiving) DCE to signal the center of each bit sent by the remote (transmitting) DCE.

Pin 24. Transmitter Signal Timing Element (to DCE). Circuit DA (ITU-T: 113). If the DCE does not supply its own clocking signal (to determine the center of each bit), the signal must be supplied on this circuit by the DTE. The on/off transition of the timing signal indicates the center of each bit.

Pin 15. Transmitter Signal Timing Element (from DCE). Circuit DB (ITU-T: 114). If the DTE does not supply its own clocking signal, the signal is supplied by the DCE on this circuit. The DTE supplies a bit between each timing signal.

Pins 11, 18, 25. Unassigned.

8.4.6 9-Pin Connectors

From the above discussion on RS-232-C, it should be clear that the majority of the pins on the standard 25-pin connector (also known as a DB-25 connector) usually can be safely ignored. For that reason, many manufacturers now produce 9-pin (or DB-9) RS-232-C connectors. The assignment of pins on an RS-232-C 9-pin connector is shown in Table 8.3 (note that there is *not* a one-to-one correspondence with the 25-pin connector).

8.4.7 The Null-Modem

In theory, two DTEs cannot communicate unless they are connected via a pair of DCEs, because the transmit and receive pins connect directly. However, in Chapter 4 it was shown that if two DTEs are spaced less than 50

Table 8.3 RS-232 9-pin signal names.

9-Pin Connector	25-pin Equivalent	Signal Name
1	8	Carrier Detect
2	2	Transmit Data
3	3	Receive Data
4	20	Data Terminal Ready
5	7	Signal Ground
6	6	Data Set Ready
7	4	Request to Send
8	5	Clear to Send
9	22	Ring Indicator

feet (15 meters) apart, a special null-modem cable will permit them to communicate. The null-modem cable crosses pins 2 (transmit) and 3 (receive), thereby ensuring that whatever is sent by one DTE will be received by the other. The minimum null-modem wiring is shown in Figure 8.21.

In many cases, application software is written to support more than simply the transmit and receive connections. For example, an application written for half-duplex communications will probably not work with the null-modem cable in Figure 8.21 because it expects signals such as Clear to Send to be set. A complete null-modem cable wiring (supporting both half- and full-duplex communications) is shown in Figure 8.22 (the diagram assumes an RS-232-C 25-pin connector).

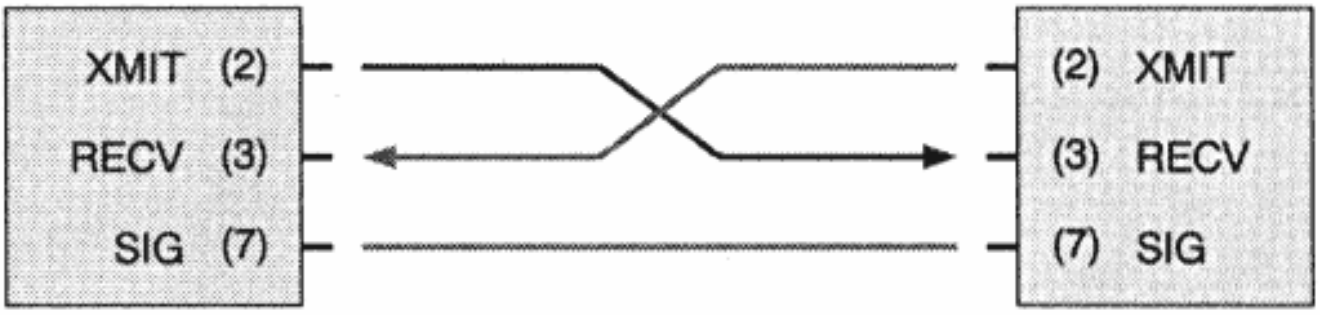

Figure 8.21
The minimum null-modem wiring.

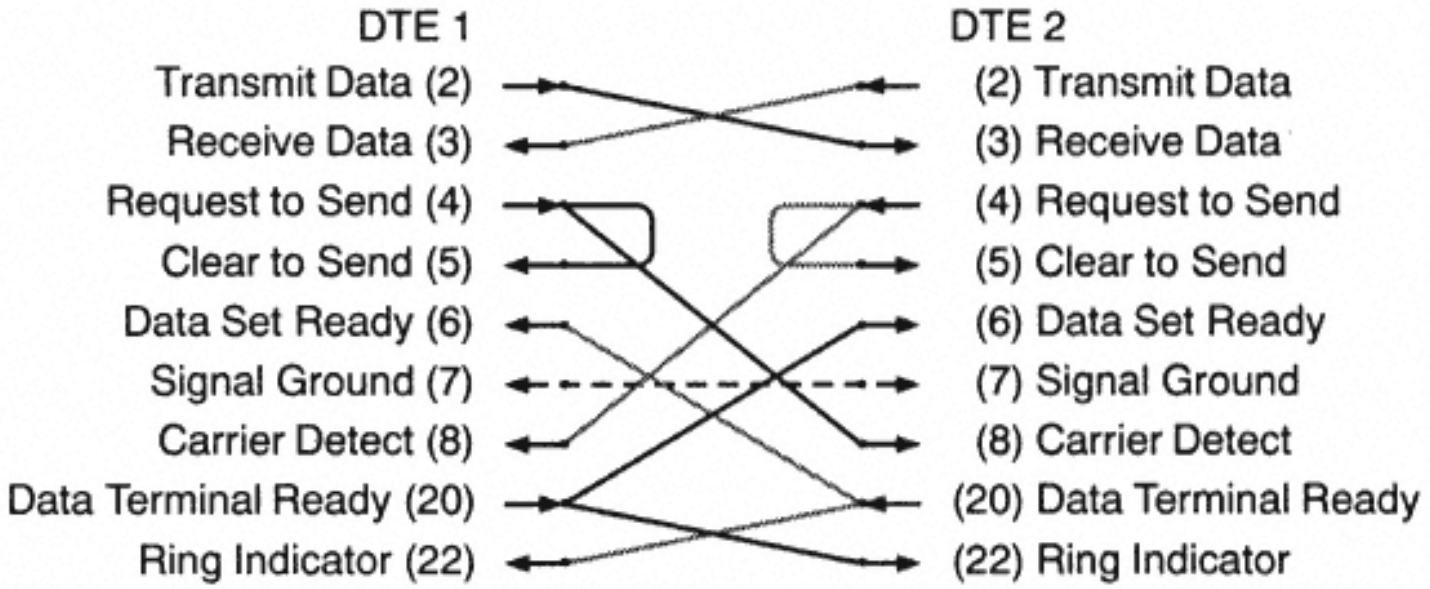

Figure 8.22
A complete null-modem wiring.

The results of the above connections cause the following to occur:

Local DTR *on* turns remote RI *on* and remote DSR *on*.

Local RTS *on* turns local CTS *on* and remote CD *on*.

Local TxD *on* turns remote RxD *on*.

8.4.8
Limited-Distance Modems

In certain situations, the distance between the two DTEs does not warrant the installation of a separate telephone circuit and modem, but at the same time RS-232-C signals are not strong enough to reach the remote DTE. In these situations, a limited-distance modem can be employed.

Limited-distance modems draw power from the DTE's RS-232-C port to help boost the RS-232-C signal (for example, power can be taken from pin 4 (RTS) or 20 (DTR)). The maximum separation between the two DTEs depends upon the manufacturer of the limited-distance modem, but typically cannot exceed about 2.5 kilometers.

8.4.9
RS-422-A, RS-423-A, and RS-449

In Chapter 4, several other connection standards were discussed, notably RS-422-A, RS-423-A, and RS-449. These three standards, like RS-232-C, define the physical and electrical standards of a DTE-DCE interconnection. Table 8.4 shows the RS-449 with its corresponding RS-232-C and ITU-T V.24 signals.

Although the Electrical Industries Association intends to replace RS-232-C with RS-422-A, RS-423-A, and RS-449, the popularity and widespread use of RS-232-C means that the replacement process will probably be a slow one. In addition, the trend in DTEDCE interconnection is towards *fewer* wires rather than more, meaning that these new standards may already be outdated.

8.5
High-Speed (Synchronous) Modems

Until recently, the vast majority of PCs connected to the telephone network used low-speed modems that offered data rates from 200 bps to 1200 bps. Although higher-speed synchronous modems were available (offering speeds up to 9600 bps), the costs, often 20 to 30 times the cost of a 1200-bps asynchronous modem, made them prohibitively expensive for a single PC.

Today however, the cost has dropped to the point where a high-speed synchronous modem is affordable to most PC owners. Furthermore, through techniques such as compression and error control, the effective data rate can exceed 56,000 bps. The components of a high-speed synchronous modem are shown in Figure 8.23.

Table 8.4 RS-449 and corresponding RS-232 and ITU-T V.24 signals.

	RS-449		RS-232-C		ITU-T V.24
SG	Signal Ground	AB	Signal Ground	102	Signal Ground
SC	Send Common			102a	DTE Common
RC	Receive Common			102b	DCE Common
IS	Terminal in Service				
IC	Incoming Call	CE	Ring Indicator	125	Calling Indicator
TR	Terminal Ready	DC	Data Terminal Ready	108/2	Data Terminal Ready
DM	Data Mode	CC	Data Set Ready	107	Data Set Ready
SD	Send Data	BA	Transmitted Data	103	Transmitted Data
RD	Receive Data	BB	Received Data	104	Received Data
TT	Terminal Timing	DA	Transmitter Signal Element Timing (DTE)	113	Transmitter Signal Element Timing (DTE)
ST	Send Timing	DB	Transmitter Signal Element Timing (DCE)	114	Transmitter Signal Element Timing (DCE)
RT	Receive Timing	DD	Receive Signal Element Timing	115	Receiver Signal Element Timing (DCE)
RS	Request to Send	CA	Request to Send	105	Request to Send
CS	Clear to Send	CB	Clear to Send	106	Ready for Sending
RR	Receiver Ready	CF	Received Line Signal Detector	109	Data Channel Received Line Signal Detector
SQ	Signal Quality	CG	Signal Quality Detector	110	Data Signal Quality Detector
NS	New Signal				
SF	Select Frequency			126	Select Transmit Frequency
SR	Signaling Rate	CH	Data Signal Rate Selector (DTE)	111	Data Signaling Rate Selector (DTE)
SI	Signaling Rate Indicator	Cl	Data Signal Rate Selector (DCE)	112	Data Signaling Rate Selector (DCE)
SSD	Secondary Send Data	SBA	Secondary Transmitted Data	118	Transmitted Backward Channel Data
SRD	Secondary Receive Data	SBB	Secondary Received Data	119	Received Backward Channel Data
SRS	Secondary Request to Send	SCA	Secondary Request to Send	120	Transmit Backward Channel Line Signal
SCS	Secondary Clear to Send	SCB	Secondary Clear to Send	121	Backward Channel Ready
SRR	Secondary Receiver Ready	SCF	Secondary Received Line Signal Detector	122	Backward Channel Received Line Signal
LL	Local Loopback			141	Local Loopback
RL	Remote Loopback			140	Remote Loopback
TM	Test Mode			142	Test Indicator
SS	Select Standby			116	Select Standby
SB	Standby Indicator			117	Standby Indicator

8.5.1
Interchange Circuits

The interchange circuits in a typical high-speed modem permit the exchange of data with a DTE. The DTEDCE connection is typically RS-232 or V.24.

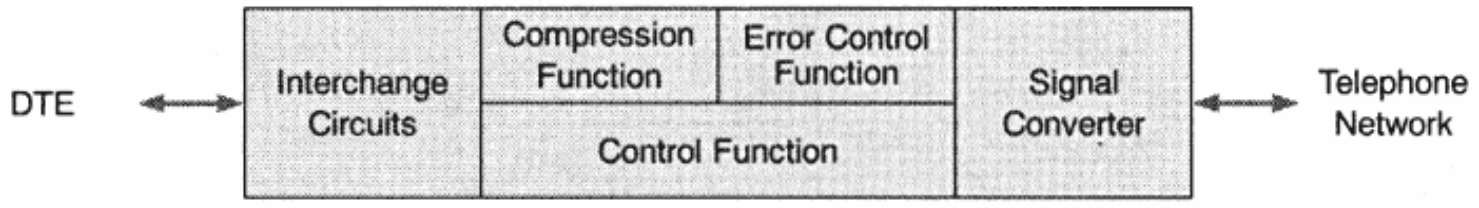

Figure 8.23
A schematic diagram of a high-speed modem.

8.5.2
Control Function

High-speed modems differ from their low-speed counterparts in several important areas, including the ability to negotiate various options as well as to turn functions off and on. These operations are internal to the DCE and transparent to the DTE. For example, ITU-T recommendation V.42 requires that a control function be responsible for the following:

Communicating with the remote DCE to negotiate the channel speed (14.4 kbps, 9.6 kbps, 7.2 kbps, or 4.8 kbps), error correction (on or off), data compression (on or off), and frame format (bit or byte oriented).

Handling all communications with the local DTE via the V.24 interface, checking for flow control signals (X-ON and X-OFF), indicating breaks, and converting asynchronous (i.e., start-stop) characters to synchronous and vice versa.

Changing parameters dynamically as conditions warrant; for example, error correction, data compression, and channel speed can change over the course of a transmission.

8.5.3
Compression Function

An examination of any text will reveal strings or groupings of characters that are identical; for example, in English, the words *the* and *and*. as well as words ending in *ion* and *th*. are extremely common. By encoding these common strings in shorter code words, the number of bits required for a transmission decreases. As an example, if the string *the* were transmitted synchronously (i.e., without start and stop bits), it would occupy 24 bits, at 8 bits per character. However, if common combinations of characters could be encoded into, say, 11 bits, then each occurrence of *the* would occupy only $\frac{11}{24}$ or about 46 percent of its original size.

To simplify the encoding process, one typically finds that all strings, including single characters, share a single code-word size. Thus, if the encoding algorithm encodes into 11-bit code words, a total of 211, or 2048 code words, is possible. Of these 2048 code words, a number will be "reserved" for the ASCII (or EBCDIC) character sets (this means, of course, that 8-bit characters are encoded into 11-bit code wordsa waste of 3 bits per code word). Each code word is stored in a dictionary of strings, known as the encoding dictionary.

To retrieve the original data, it is necessary to decode the code words, which can be done using table lookup. For 11-bit code words, a table of 2048 entries can be used, with each one corresponding to the equivalent decoded string. This table is known as the decoding dictionary.

The encodedecode algorithm described above is essentially the same as ITU-T Recommendation V.42*bis*. That is, it scans a stream of characters and builds the corresponding code word as each character is encountered. When the next character read results in a string that does not exist in the encoding dictionary, the code word for the string up to (but not including) the last character read is transmitted to the remote DCE.

The V42*bis* encoding dictionary can be represented as a set of trees; as an example, the tree with root *S*, and representing the strings *S*, *SA*, *SAD*. and *SAT*. is shown in Figure 8.24. Each node of the tree is associated with a unique code word (shown beside each letter).

When a new character is encountered, the root of the tree is found; then the tree is traversed as each subsequent character is read. When the end-of-string is detected (i.e., a leaf node) or when there is no corresponding letter in the subtree, the code word has been found. If, for example, the letters *SAT* were supplied to the V.42*bis* algorithm with the above tree, the code word generated would be *291*.

With a limited number of possible code words (for example, 11 bits permit only 2048 code words), a problem arises: what is the initial value of the dictionary? One solution is to have a set of what the manufacturer believes to be the most common "strings" placed in the dictionary beforehand. If this (static) solution is adopted, the encoding algorithm will work best for those strings supplied by the manufacturer, but in other cases it may be more costly to use encoding (because of the additional bits associated with each character).

In V.42*bis*. the dictionary is allowed to evolve dynamically over time. For example, if the string *SAN* is supplied to the encoding algorithm, the characters *SA* are encoded as code word *307*; then the character *N* is added to the dictionary using the next available (or oldest) keyword, say *452*. as shown in Figure 8.25. Processing resumes using *N* as the first character in a new dictionary search (*not* as part of the current search).

If a compression algorithm is to work, it is necessary for both the encod-

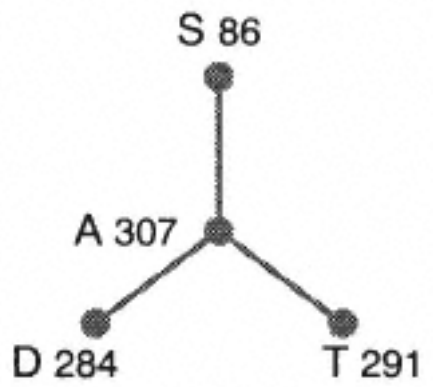

Figure 8.24
An example of a V.42*bis*
encoding dictionary entry.

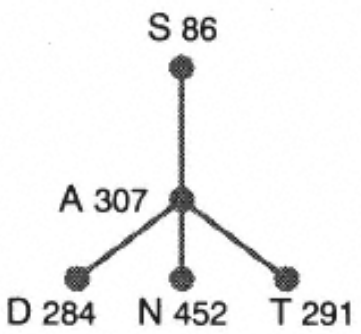

Figure 8.25
Adding a new entry to
a dictionary tree.

ing and decoding algorithms to remain in step. A brief examination of the last example shows that in V.42*bis* this is not possible because the decoding algorithm will always be one character behind the encoding algorithm. In the example, the encoder recognizes that *SAN* is not in the dictionary and therefore sends the code word for *SA*307and makes *N* part of the *S* tree.

The decoder receives *307*. but has no indication as to the next character until the following code word is decodedin this case *N*. the first character of the next code word. By applying the same algorithm used by the encoder to update the encoding dictionary, the decoder's dictionary will appear identical to that of the encoder, albeit a character behind. Figure 8.26 shows the relationship between the encode and decode algorithms and the dictionary algorithm.

In V42*bis*. the control functions in both DCEs negotiate the number of bits in a code word when the call is initialized. Over time, if more code words are required, the two DCEs can negotiate an increase in code word size (by increasing the number of bits in the code word, the number of possible code words increases). Each DCE has two dictionaries, one for encoding and one for decoding; their values need not be the same, since the data transmitted by a DCE may not be the same as that received.

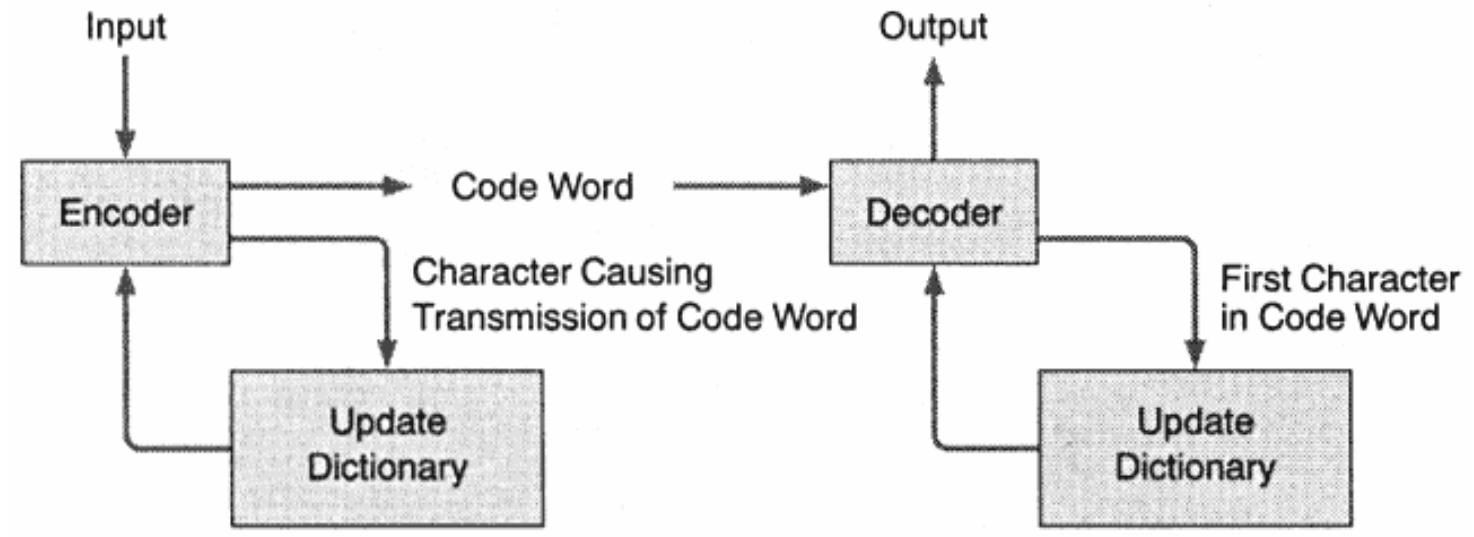

Figure 8.26
The relationship between the dictionary, the encoder, and the decoder.

8.5.4
Error Control Function

In Chapter 6, it was demonstrated that synchronous communications offer several advantages over asynchronous communicationsnotably, the reduction of transmission overheads as well as improved error detection and correction facilities. Furthermore, bit-oriented synchronous protocols such as SDLC permit the transmission of streams of bits (not necessarily in multiples of eight).

Both of these capabilities are, in part, the basis of the high-speed modem error-control functions. First, by incorporating a reliable protocol (such as SDLC) within the modem, the application need not be concerned with activities such as the retransmission of bad frames. Second, bit-oriented protocols are more applicable to handling compressed data, since the code words generated by V.42*bis* compression need not be a multiple of eight.

Many high-speed modems now support ITU-T Recommendation V.42, which defines a protocol for error detection and subsequent correction for use between DCEs. The protocol, a subset of the ITU-T HDLC (which, in turn, is a variant on SDLC), is referred to as LAP-M or Link Access Procedure for Modems.

The basic structure of the LAP-M frame is identical to that of the SDLC frame (i.e., each frame is delimited by a pair of flag bytes, `01111110`, and consists of an address field, a control field, an information field, and a 16- or 32-bit FCS). The 16-bit FCS is CRC-CCITT, while the 32-bit FCS is CRC-32; CRC-32 is recommended when using compression.

LAP-M has three frame formats, notably *I* (information), *S* (supervisory), and *U* (unnumbered). The control field for the *I* and *S* frames differs slightly from SDLC in that the send and receive sequence numbers, *Ns* and *Nr*. are 7 bits rather than 3 bits in length.

An alternative to the bit-oriented LAP-M frame is a byte-oriented frame that conforms to the Bisync protocol discussed in Chapter 6. The frame is sent transparently; each character is sent asynchronously (i.e., with start and stop bits).

8.5.5
Signal Conversion

In any communication, both the transmitter and the receiver must remain in step to ensure that the data transmitted is the same as the data received. In transmission of data between two asynchronous DCEs, the communication is much the same as between two UARTs, using start and stop bits while limiting the number of bits transmitted.

In data transmission between two synchronous DCEs, there is no limit to the number of bits that can be sent in theory. To avoid the loss of data through, for example, clock drift, it is necessary for both DCEs to remain synchronized; this can be achieved by incorporating highly accurate clocks

within each DCE (one of the reasons for the higher costs associated with synchronous modems) and keeping them synchronized for the duration of the transmission.

Section 8.3.2 shows that a dibit (two bits per baud) could be sent by phase shifting with four different phase angles. This can be increased to three bits per baud (the tribit); for example, ITU-T Recommendation V.27 supports 4800 bps on a 1600-baud channel. The phase encoding for V.27 is given in Table 8.5.

To keep the two DCEs synchronized, it is necessary to ensure that there are sufficient phase changes in the data stream for the receiving DCE to correct its clock (this is similar to Manchester encoding, discussed in Chapter 6). If the data supplied by the transmitting DTE does not change over a long period of time (for example, a stream of 0s), there will be no change in the phase angle, meaning that the DCEs' clocks can lose synchronization. Thus, the problem is how to introduce phase changes into the data stream.

The solution adopted for synchronous modems is to incorporate a scrambler into the data stream before the bits are encoded. The scrambler is usually designed as a feedback register (similar to the CRC shift register described in Chapter 6) that takes a series of bits and produces a new bit pattern. The bit patterns generated are intended to ensure that all possible phase angles are equally likely to occur. For example, the generating polynomial for the V.27 scrambler is $1 + x6 + x7$. The receiving DTE has a corresponding descrambler that takes the decoded bit stream and extracts the correct bit values.

Speeds higher than 4800 bps are also possible, but these typically do not rely on phase shifting alone; instead, they incorporate phase shifting and amplitude modulation. ITU-T Recommendation V.29, for example, supporting 4 bits per baud and operating at 2400 baud, produces 9600 bps. The leftmost data bit (designated Q1) determines the amplitude (see Table 8.6),

Table 8.5 Phase encoding for V.27.

Tribit	Phase Differential
001	0
000	45
010	90
011	135
111	180
110	225
100	270
101	315

Table 8.6 Amplitude, determined from phase and Q1.

Absolute Phase	Q1	Signal Amplitude
0°, 90°, 180º, 270°	0	3
	1	5
45°, 135°, 225°, 315°	0	$\sqrt{2}$
	1	$3\sqrt{2}$

Table 8.7 Phase shift is determined by Q2, Q3, and Q4.

Q2	Q3	Q4	Phase
0	0	1	0º
0	0	0	45°
0	1	0	90°
0	1	1	135°
1	1	1	180°
1	1	0	225º
1	0	0	270°
1	0	1	315°

while the remaining three bits (designated Q2, Q3, and Q4) determine the phase change (see Table 8.7). The constellation pattern for V.29 at 9600 bps is shown in Figure 8.27.

Speeds as high as 14,400 bps (14.4 kbps) are achievable using ITU-T Recommendation V.32*bis*. in which groups of six data bits (designated Q1, Q2, Q3, Q4, Q5, and Q6) are encoded using a trellis encoder. The product

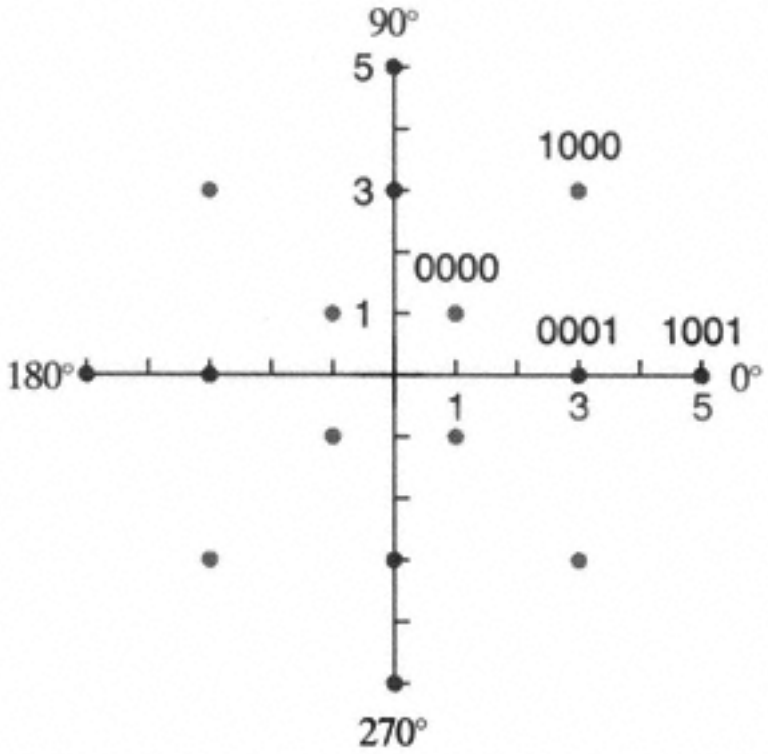

Figure 8.27
Constellation pattern for ITU-T
Recommendation V.29 at 9600 bps.

of the encoding is *seven* bits (bits Y0, Y1, and Y2, from data bits Q1 and Q2, as well as the remaining data bits Q3, Q4, Q5, and Q6). These bits map into a constellation pattern of 128 possible values that is sent to the remote DCE, where it is decoded and the data bits are restored.

V.32*bis* supports lower line speeds by reducing the number of data bits: at 12,000 bps (12 kbps) Q6 is not used; nor are Q6 and Q5 at 9600 bps; and at 4800 bps Q6, Q5, and Q4 are ignored. Bits Q1 and Q2 are always passed through the encoder; bits Y0, Y1, and Y2 are always transmitted. At 9600 bps, for example, five bits are required to encode the four data bits.

8.5.6 Facsimile

The facsimile, or more commonly, the fax machine, allows the transmission of documents across the telephone system. ITU has developed a number of facsimile standards; the two in most widespread use are Group 3 (for analog public telephone networks) and Group 4 (for digital public telephone networks). Although there are many fax machines on the market, the basic operations are the same: pages of the document are digitized and then transmitted.

A document page is digitized into a series of bits referred to as picture elements or pels, which are either 0 (white) or 1 (black). The horizontal density is 8.05 pels per millimeter, with 3.85 or 7.7 lines per millimeter (depending upon the quality selected). A page of, say, $8\frac{1}{2}$ by 11 inches (roughly 210 by 270 millimeters), when digitized, produces between 1.8 and 3.5 million pels. Given a 9600-bps asynchronous channel, the time to send a page this size is in the range of 3.8 to 7.6 minutes.

To reduce this time, the digitized pages are compressed, and sequences of white or black pels are recognized and represented as a unique bit pattern. Unlike V.42, the code directory is not dynamic. Specific sequences of bits have unique code words that do not change over time; thus, a string of 60 white pels is encoded as `01001011`, whereas `000000101100` represents a string of 60 black pels. Sequences of pels that are longer than multiples of 64 (64, 128, 192, and so on) are encoded into makeup codes (the nearest multiple of 64) and termination codes (the difference between the sequence length and the makeup code). For example, a sequence of 68 white pels is sent as the makeup code for 64 white pels plus the termination code for 4 white pels. With compression, a Group 3 facsimile can transmit a page in less than a minute and a Group 4 can transmit a page in several seconds.

Line synchronization is maintained by assuming that each line starts with a single white pel; this ensures that the transmitter and receiver work from a common starting point. If a Group 3 facsimile loses bit synchronization, it must scan the incoming bit stream for a special end-of-line sequence, discarding the bits up to that point. In a Group 4 facsimile, the receiving fax machine can request a retransmission of the data in error using a variant of HDLC.

8.6 PC Modem Control

The 8250 UART supplied with the PC supports a limited number of RS-232-C modem functions, specifically:

The ability to control a number of modem conditions (through the Modem Control Register).

The ability to determine the modem's status (from the Modem Status Register).

The UART can be programmed to interrupt the PC when modem status changes occur.

8.6.1 The Modem Control Register

Previous chapters demonstrated a number of uses for the Modem Control Register (port `0x3FC` or `0x2FC`), including loopback and enabling of UART interrupts. The Modem Control Register can also be used by the PC to send DTR (Data Terminal Ready) and RTS (Request to Send) signals to the UART. Its structure is shown in Figure 8.28, and its bits are defined as follows:

DTR (Data Terminal Ready). When set, the modem is signaled that the PC has connected to the modem and is able to send and receive data.

RTS (Request to Send). The PC is requesting access to the channel. This is often ignored in full-duplex modems, but is necessary for half-duplex configurations.

Out 1. Not used.

Out 2. Must be set to enable UART interrupts.

Loop. Set to enter loopback mode (see Chapter 5).

Bits 5, 6, 7. Permanently 0.

For example, to signal the modem that the PC is connected to the channel and has data to send, one could write:

```
#define MCR      0x3FC
#define DTR      0x01
```

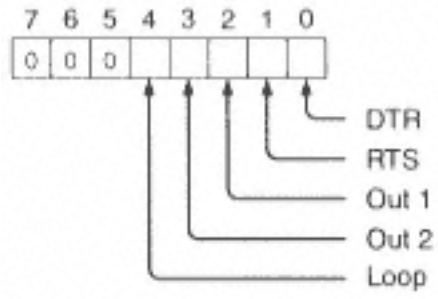

Figure 8.28
The 8250 Modem Control Register.

```
#define RTS      0x02
#define OUT2     0x08


outportb(MCR, DTR+RTS+OUT2)
```

To clear the modem, one could use the following (note that `OUT2` must be written to the Modem Control Register to ensure that all serial port interrupts will be allowed to continue):

```
outportb(MCR, OUT2);
```

8.6.2
The Modem Status Register

The Modem Status Register (port address `0x3FE` or `0x2FE`) can be used to determine a limited number of modem status indications. Specifically, it indicates a change in state of any of four RS-232-C connections as well as the current value of connection. The Modem Status Register is shown in Figure 8.29. Its bits are defined as follows:

ΔCTS. There has been a change in the Clear to Send signal since the last time the Modem Status Register was read. The new value of CTS can be obtained by examining bit `0x10` of the Modem Status Register.

ΔDSR. There has been a change in the Data Set Ready signal since the last time the Modem Status Register was read. The new value of DSR can be obtained by examining bit `0x20` of the register.

ΔRI. There has been a change in the Ring Indicator signal since the last time the Modem Status Register was read. The new value of RI can be obtained by examining bit `0x40` of the register. This bit is set each time the telephone rings and stops ringing.

ΔCD. There has been a change in the Carrier Detect signal since the last time the Modem Status Register was read. The new value of CD can be obtained by examining bit `0x80` of the register.

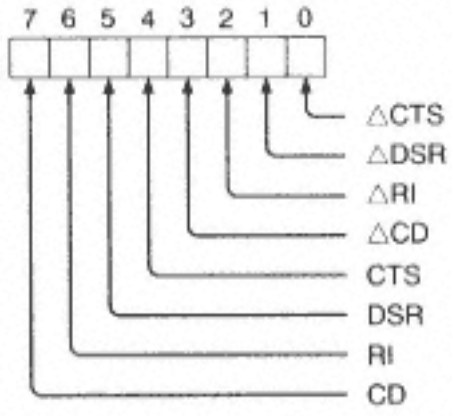

Figure 8.29
The 8250 Modem Status Register.

CTS. The value of the Clear to Send signal. In full-duplex applications, CTS is usually permanently set; in half-duplex applications its value depends upon whether the DTE is attempting to transmit.

DSR. The value of the Data Set Ready signal. If the modem is functioning properly, this signal will be on; otherwise, it is cleared. For data transfer to proceed, DSR must be on.

RI. The value of the Ring Indicator signal. If a ringing tone is detected, RI is set *on*; otherwise it is set *off*. Ring Indicator refers solely to answer-only modems. Once the call is established, RI should not be generated.

CD. The value of the Carrier Detect (Received Line Signal Detector) signal. If a carrier is present, data transfer can take place, subject to the rules of the protocol.

Depending upon the signals available from the modem, a data transfer should wait for CTS, DSR, and CD. For example, to wait for modem signals from serial port 2, one could write:

```
status = inportb(SP2 + 4);

 if ((status & (CTS+DSR+CD)) == (CTS+DSR+CD))
      /* Data transfer can now proceed */
```

8.6.3 Modem Interrupts

Modem status changes can be made to cause interrupts by setting the UART's Interrupt Enable Register to `0x08`. Then, whenever a change in the modem status occurs, an interrupt will occur and the UART's Interrupt Identification Register will contain a value of `0x00` (indicating a modem status change). The status of the modem can be obtained by reading the Modem Status Register:

```
while ((iir = inportb(address + 2)) != 1)
{
   switch (iir)
   {
   case 6: /* Line status */
      . . .
   case 4: /* Data available */
      . . .
   case 2: /* Transmit done */
      . . .
   case 0: /* Modem status change */
      status = inportb(address + 6);
      /* Process status */
   }
}
```

8.7 Modem Testing

Although the addition of modems and the telephone network extends the distance over which a DTE can communicate, it also means that more things can go wrong. Fortunately, most modems support a number of tests that can assist in tracking down communication errors.

8.7.1 Self-Tests

Some modems can perform self-tests in which a specific binary test pattern is generated and then modulated. The modulated signal is looped back and demodulated. If the demodulated pattern is the same as the original test pattern, the self-test is successful. Ideally, both the originate and answer frequencies are tested. If an error is detected, it is usually signaled through an error indication light on the front panel of the modem.

During a self-test, the connections to the DTE and the telephone network are disconnected; the internal configuration of the modem self-test is shown in Figure 8.30.

8.7.2 The Analog Loopback Test

An analog loopback test is designed to check the operation of the local DTE. Signals transmitted by the DTE enter the modem, where they are modulated, demodulated, and then returned to the DTE. To allow modem software to be tested, carrier detect is set high. The test is successful if the data transmitted is equal to the data received.

In an analog loopback test, the DCE is isolated from the telephone network; Figure 8.31 illustrates the configuration.

8.7.3 The Digital Loopback Test

The digital loopback test is intended to test local demodulation and modulation of remote data; the local DTE does *not* participate. The remote DTE/DCE transmits modulated data across the telephone network, and the local DCE demodulates it and loops it back through the modulation circuitry.

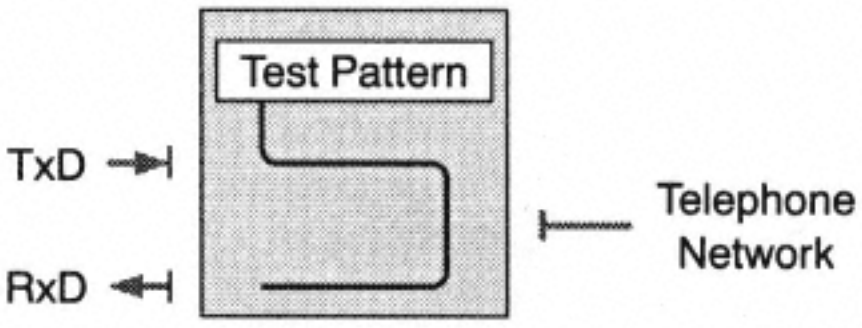

Figure 8.30
Modem self-test.

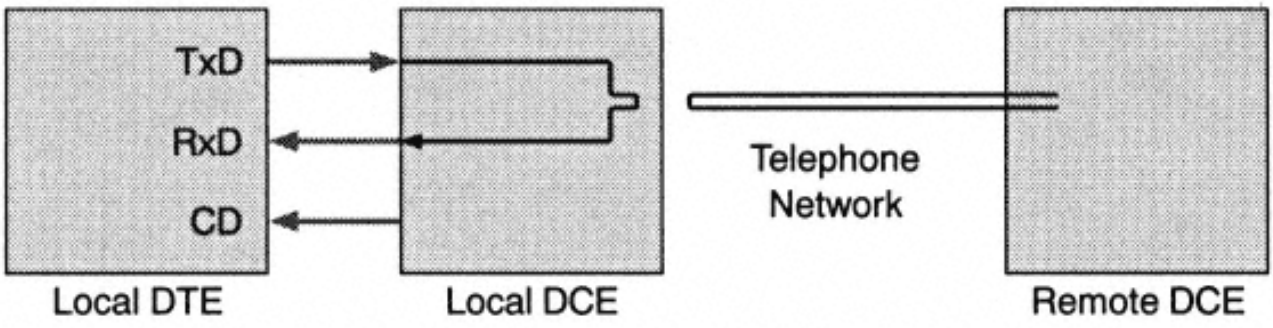

Figure 8.31
Analog loopback test.

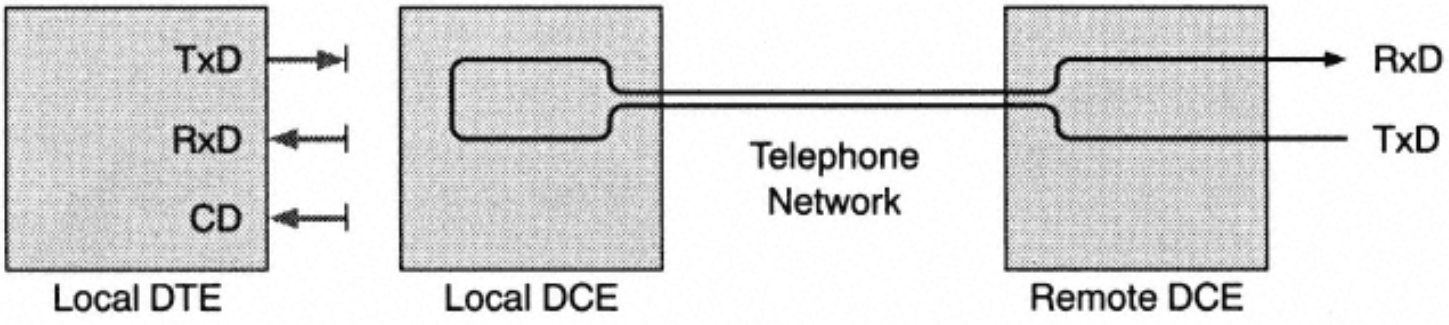

Figure 8.32
Digital loopback test.

The configuration of the digital loopback test is shown in Figure 8.32. The test is successful if the data received by the remote DTE is the same as that sent by it.

A variation on the digital loopback is the modified digital loopback, in which the local transmitter is looped back at the modem (see Figure 8.33).

8.7.4
The Remote Loopback Test

The remote loopback test permits testing of both modems and the telephone network. The configuration is shown in Figure 8.34. Signals are transmitted across the telephone network and looped back at the remote DCE, which enters digital loopback test mode. The tests are successful if, at the local DTE, the data sent is the same as that received.

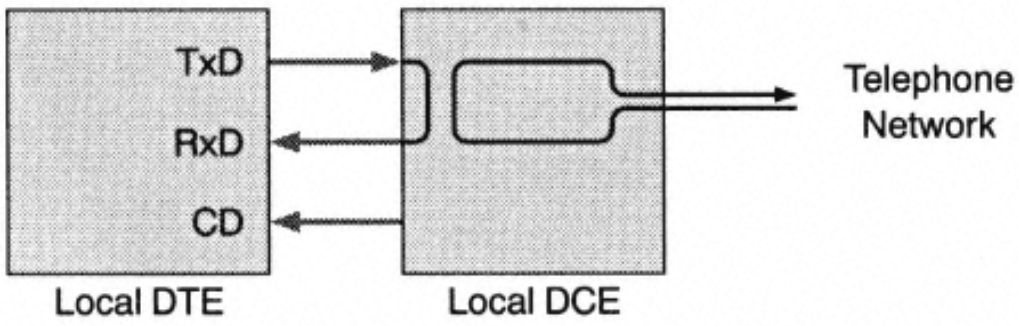

Figure 8.33
Modified digital loopback test.

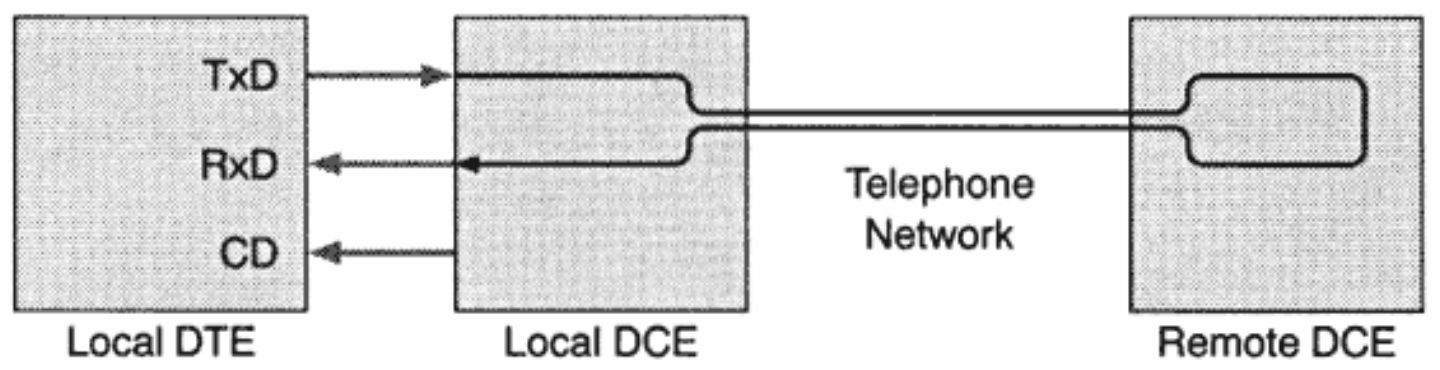

Figure 8.34
The remote loopback test configuration.

8.8
Modem Variations

A number of interesting extensions to the modem are available commercially. One group supports textual commands known as Hayes' commands; another is referred to as automatic calling units, or ACUs.

8.8.1
Hayes' Commands

A limitation associated with programming the PC's UART to handle a modem directly is that not all manufacturers conform to the RS-232 (V.24) signals. In certain cases, software written (and tested) for one type of modem will not function for a second "identical" modem from a different manufacturer. To overcome these limitations, many modems support Hayes' commands, a set of textual strings that are recognized and interpreted by the modem, thereby eliminating the need for special software to be written for each modem on the market.

To support Hayes' commands, a modem requires a front-end that interprets the modem commands issued by the application software, converting them into the signals required by the modem. To simplify the design, most front-ends appear as a UART and commands are issued as text strings. To drop carrier, for example, the application software generates the string `ATCO`. Similarly, when a change in line status is detected by the modem, it is returned to the application as a string; thus, detecting a ring causes the front-end to return the string `RING`. All commands and responses pass through the UART's transmit and receive buffers.

When first initialized, the front-end is in the "command" state, interpreting any data supplied to it. Once a connection is established with another modem, the front-end enters the "on-line" state, ignoring all data and passing it directly to the modem. The front-end can be taken out of the on-line state and placed in the command state by sending a sequence of three "+"s. It is the responsibility of the application software to know the state of the front-end.

Some of the more widely used Hayes' commands are

`DTx`. Dial telephone number *x* using tones (as opposed to pulses). If a connection is made, the modem enters the on-line state.

`&Zn = x`. Assign a telephone number, *x*, to an internal telephone register, *n*.
`S = n`. Dial the telephone number stored in register *n*.
`H`. Hang up the call (the telephone goes on-hook).
`O`. Change from command state to on-line state.

Note that all of the above commands must be prefixed by `AT` (where `AT` signals *AT*tention) in order to be recognized by the modem. If the command is recognized and performed, the modem responds `OK`.

8.8.2
Automatic Calling Units

An automatic calling unit is a hardware/software package that can place an outgoing call to a remote DTE.

One common application of the ACU is to place data transfer calls late at night to take advantage of low-cost telephone rates. Another is to permit the source of the incoming call to be verified. In the latter situation, the DTE with the ACU is called and obtains an identification associated with the calling DTE. It can then contact that DTE to verify that this was the source of the call. If it was, the data transfer can proceed.

The EIA standard RS-366 provides specifications for automatic calling units.

8.9
Commkit External Modem Software

Commkit is supplied with software that demonstrates how the PC and an external modem can be made to interact. The software, `external.c`, extends the point-to-point telephone utility developed in Chapter 3 to handle modems. `External.c` supports both originating and answering modems.

8.9.1
Hardware Requirements

Two test configurations are possible. For example, as illustrated in Figure 8.35, if two PCs, two external modems, and two telephones are available, all of

Figure 8.35
Hardware configuration if two modems are available.

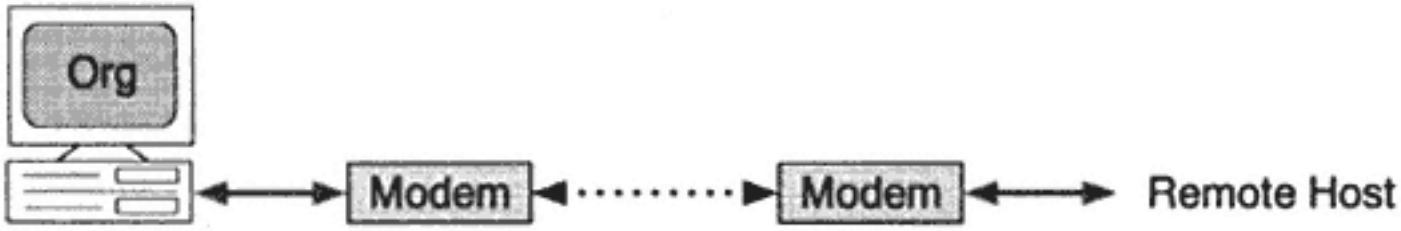

Figure 8.36
Hardware configuration if only one modem is available.

the testing can take place using the modem software (Org is the originating station; Ans the answering station).

However, testing can still take place with a single modem if a remote host with dial-in facilities is available. In this case, an outgoing call is placed to the remote host; output from the remote host will appear on the lower half of the screen, while the local input will appear on the upper half. The hardware configuration is shown in Figure 8.36.

Finally, regardless of the configuration, the PC and the modem must be connected by a straight-through rather than a null-modem cable.

8.9.2
Design

The original point-to-point software will not support a modem. Any modem status changes detected by the `commkit` module and passed to `low_level()` with a `code` value of `MODEMSTATUS` are simply ignored. Modifications must be made if modem status changes are to be recognized and acted upon.

The modem status changes that are recognized by the 8250 UART are limited to

Clear to Send. The modem is ready for transmission.

Data Set Ready. The modem is in data mode.

Ring Indicator. The modem is receiving a ring from the telephone line.

Received Line Signal Detect. The modem is receiving carrier (also referred to as Carrier Detect).

The major difference between the answering modem and the originating modem is whether Ring Indicator is detected. (Ring Indicator should be detected only by the answering modem.)

Once carrier is detected (by the answering modem asserting carrier or the user at the originating modem pressing the Data/Talk button), the modem should enter the Data Set Ready state and signal both `DSR` and `RLSD` to the UART. Upon detection of `DSR` and `RLSD`, the DTE can respond with Data Terminal Ready and Request to Send (since these are the only two signals supported by the UART). At this point, the DTE is connected to the DCE and communications can commence.

Communications will continue until either end drops the carrier or until the detection of a CTRL-C.

8.9.3
Implementation

Modifying the Commkit point-to-point software to support modems essentially entails adding modem control software to monitor the state of the modem before and during the exchange of information between DTEs. A two-state finite-state machine is defined in Figure 8.37; it controls both an originating and an answering modem (`mcr` is the value of the Modem Control Register, and `msr` is the value of the Modem Status Register).

When a modem status change occurs, control is passed to `low_level()` with `device` indicating the serial port, a `code` value of `MODEMSTATUS`, and the value of the Modem Status Register in `data`. Although it would be a simple matter to forward the value of `data` to a process and have the process check the modem status, the low-level Commkit modem software interprets the modem status and forwards a code to the foreground process. Similarly, the foreground process does not access the Modem Control Register; instead, a message is sent to the low-level modem software indicating the required actions.

To support this design, the following changes are necessary:

1. A mechanism whereby the foreground process can send messages to the serial port for transmission or for controlling the modem. A number of solutions are possible, such as sending a 2-byte message to the modem serial port: one byte containing the data, the other containing a code indicating whether the byte is for transmission or to control the modem. The solution adopted involves creating two new process identifiers: `MODEM_DATA` and `MODEM_CHANGE`, defined as `SP1IH` and `SP2IH`, respectively. Now, all data to be transmitted is sent to `MODEM_DATA` and all modem control information is sent to `MODEM_CHANGE`.

2. `Low_level()` is modified to recognize the two process identifiers. `Code` values associated with `MODEM_DATA` are taken to mean the standard actions associated with any serial port software; the only acceptable `code` value associated with `MODEM_CHANGE` is `MSG_AVAIL`.

A `MODEM_DATA` code value of `MODEMSTATUS` *or* a message for

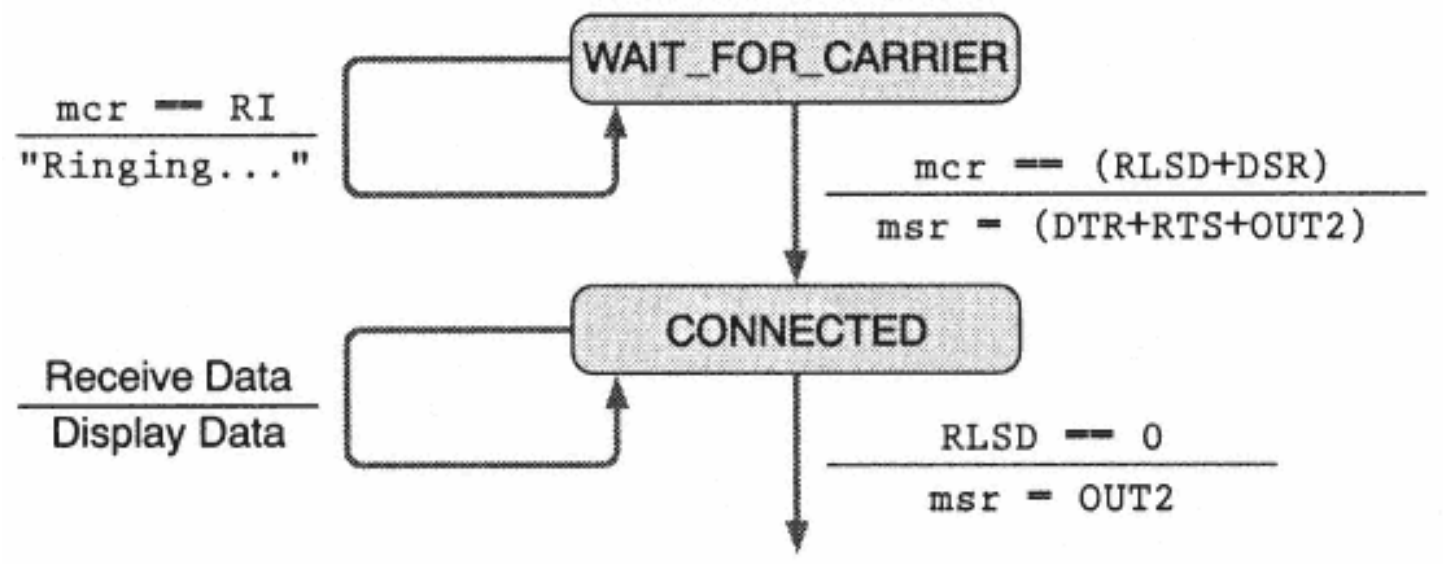

Figure 8.37
Commkit modem-control software state machine.

Table 8.8 Modem Code Values (to foreground process).

Value of msr	Modem Code
RI	RINGING
RLSD + DSR	CARRIER_AVAIL
RLSD ==0	CARRIER_LOST

MODEM_CHANGE causes control to pass to modem_change(). Modem_change() is responsible for interpreting either the value of the Modem Status Register or the message code from the foreground process:

MODEMSTATUS. Modem_change() examines the value of the Modem Status Register for a code value and, if one is found, sends the associated modem code to the foreground process. The code values are given in Table 8.8.

MODEM_CHANGE. The foreground process instructs modem_change() to change the Modem Control Register to a new value; these values are given in Table 8.9. Note that when the modem software asserts DTR and RTS, OUT2 is also asserted to ensure that the 8250 still generates interrupts. Similarly, when the call is cleared, OUT2 must still be asserted.

3. The foreground process consists of two procedures: keyboard_data() and do_modem(). Keyboard_data() accepts characters from the keyboard (via do_modem()) and displays them and forwards them to MODEM_DATA. CTRL-Cs are forwarded by keyboard_data().

Do_modem() is an implementation of the finite state machine presented at the start of this section with an additional state, DISCONNECTING. Control remains in state WAIT_FOR_CARRIER until a MODEM_CHANGE value of CARRIER_AVAIL is detected, at which point the modem control message ACCEPT_CALL is sent to MODEM_CHANGE (note that messages of type RINGING cause the string "Ringing. . ." to be displayed.

Once the call is accepted, control passes to state CONNECTED, remaining there until either the carrier is lost (message CARRIER_LOST from MODEM_CHANGE) or the user enters CTRL-C. If a CTRL-C is detected, con-

Table 8.9 New Modem Control Values (from foreground process).

Modem Code	New Value of mcr
ACCEPT_CALL	(DTR + RTS + OUT2)
CLEAR_CALL	OUT2

trol passes to state DISCONNECTING, which waits two seconds before causing the carrier to be dropped, thereby permitting any pending characters to be sent to the remote DTE.

4. No changes are required for the background process BACKGROUND_1, do_rmt_scr(); it still receives bytes to be displayed on the screen from the serial port.

8.9.4
Compiling the External Modem Software

An executable version of the external modem software, external.exe, can be created using the make utility:

```
C:\> make external.exe
```

External.c includes the header files general.h, ascii.h, and devices.h. The object file external.obj is linked with commkit.obj to produce external.exe.

8.9.5
Running the External Modem Software

Once an executable version of the modem software exists, it can be run by typing modem followed by the line speed (the modem software is written to communicate with an external modem attached to serial port 1). For example, to run the modem at 1200 bps:

```
C:\> external 1200
```

Execution begins by clearing the screen. What happens next depends upon whether the DTE is originating or answering:

Originating. The number of the remote DTE/DCE must be dialed. When the carrier is heard, the Data/Talk button must be pressed and the telephone handset replaced in the cradle; the connection is now made (which lights appear on the modem depends upon the modem). As characters are typed, the modem's Send light should flash on and off. If the parity, word size, and line speed are set correctly, a communication should take place.

Answering. The modem must be in answer mode (this might mean opening the modem and changing some settings; check the instruction manual *before* embarking on this adventure).

Have someone at the remote modem dial the number of the local modem. When the telephone rings, the word Ringing should appear on the screen. What happens next depends upon the modem. If CD and DSR are set, the connection will be made. External.c may require changes to allow the communication to take place.

Once the connection is established, whatever is typed at the local PC will be transmitted to the remote one.

8.10
Commkit Internal Modem Software

An alternative to the external modem is the *internal* modem, which is installed as a card in the PC and exhibits the properties of a UART. Most internal modems are designed to support Hayes' `AT` commands, which makes their access different from non-Hayes' modems in at least two ways:

1. The control software is not in direct communication with the UART; rather, all modem control signals are via the (textual) `AT` commands.

2. All modem status changes are interpreted by the internal modem and returned to the control software as text via the UART's data-available interrupt.

Given sufficient layers of software, these differences can be made transparent to the end-user. However, to demonstrate how the internal modem works, as well as some of its commands, the Commkit internal modem software can accept keystrokes from the keyboard and pass them directly to the modem. Therefore, all modem commands must be explicitly entered by the user; hiding them from the user is left as an exercise.

8.10.1
Design and Implementation

The Commkit internal modem software allows a user to issue commands to an internal modem as well as to send data to a remote DTE. All information received from the modem (either modem status messages or data from the remote DTE) is displayed on the PC's monitor. Message flow is shown in Figure 8.38.

`INTERNAL_IH`, shown in Figure 8.38, is the internal modem card's interrupt handler. Although most internal modem cards can be associated with

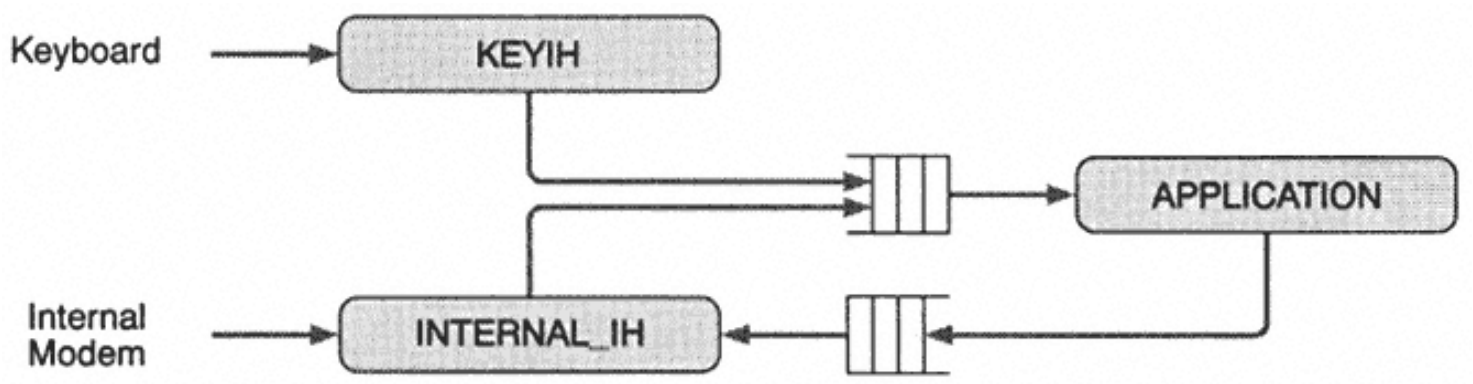

Figure 8.38
Message flow in the Commkit internal modem software.

one of a number of serial ports (including serial ports 1 and 2), a decision has been made *not* to

Require the card to map into an existing serial port.

Rewrite parts of Commkit to support a third serial port dedicated to the internal modem.

Instead, the internal modem software "cheats" by mapping all interrupts from the internal modem card onto serial port 1. This is not as difficult as it may seem; first, the interrupt vector associated with the internal modem is redirected to `spl_ih()` using `setvect()`. Thereafter, all interrupts from the internal modem are received and treated as if they were interrupts from serial port 1.

Second, the base port address of the internal modem must be available to the interrupt handler when reading the various port registers. Commkit supports a single serial-port interrupt handler, `port_handler()`, that handles interrupts from both serial ports; this is achieved by directing serial port 1 interrupts through `spl_ih()` and serial port 2 interrupts through `sp2_ih()`; see Section 3.5.1. Both `spl_ih()` and `sp2_ih()` call `port_handler()` with an identifier (`SP1IH` or `SP2IH`, respectively) indicating the source of the interrupt. `Port_handler()` uses the identifier as an index into the array `serial_addr[]; serial_addr[]` returns the base address associated with the serial port in question. At present, `serial_port[0]` has the value `0x3F8`, the base address of serial port 1; by changing this to the base address of the internal modem, all `SP1IH` interrupts will access internal modem addresses rather than the addresses of serial port 1.

The default values for the internal modem card are interrupt number 13 (IRQ 5) and base address `0x3E8`. These are defined in `internal.c` as a pair of constants:

```
#define INTERNAL_INIT   13     /* IRQ5 for this
                                  modem card */
#define INTERNAL_PORT   0x3e8  /* Base port for
                                  this modem card */
```

The initialization sequence is modified slightly (when compared to other Commkit initialization sequences) in that the interrupt vector must be established and the base address changed; two lines of code in the mainline of `internal.c` are required:

```
/* Force interrupt vector and serial port to
   modem h/w */
setvect(INTERNAL_INIT, spl_ih);
serial_addr[0] = INTERNAL_PORT;
```

8.10.2
Compiling the Internal Modem Software

If changes are made to `internal.c`, a new executable version can be created using the `make` utility:

```
C:\> make internal.exe
```

The header files `dos.h, general.h, ascii.h`, and `devices.h` are all included in `internal.c. Commkit.obj` and `srcalls.obj` are linked with `internal.obj` to make `internal.exe`.

8.10.3
Running the Internal Modem Software

Running the internal modem software involves typing `internal` followed by the desired internal modem line speed. For example, to communicate at 2400 bps, one types:

```
C:\> internal 2400
```

(Remember, the default values for the internal modem software are IRQ 5 and port address `0x3E8`.) Although some tests can be performed without connecting the modem to the telephone network, it is more instructive to have the modem connected; check the appropriate manual for details.

The internal modem software begins execution waiting for commands from the user. Whatever keystrokes are entered are interpreted as commands; for example, to access the 2400-bps Datapac node in Halifax, Nova Scotia (telephone number 453-8100), one types the following `AT` command:

```
ATDT453-8100
```

Upon detection of the carriage return following the last digit, the modem follows the specified instructions:

`AT`a command follows.

`D`take the modem off-hook and wait for a carrier (i.e., a dial tone); if none is detected, a diagnostic `NO CARRIER` is returned.

`T`if a dial tone is detected, dial the specified number using tones (as opposed to the pulses generated by a rotary dial, indicated by the command `P`).

`453-8100`the telephone number to dial.

If the connection is successful, the internal modem will respond with a `CONNECT` message (possibly followed by the line speed). Thereafter, all keystrokes pass through the internal modem directly to the remote DTE. If remote echoing is enabled, all keystrokes entered at the DTE will be echoed.

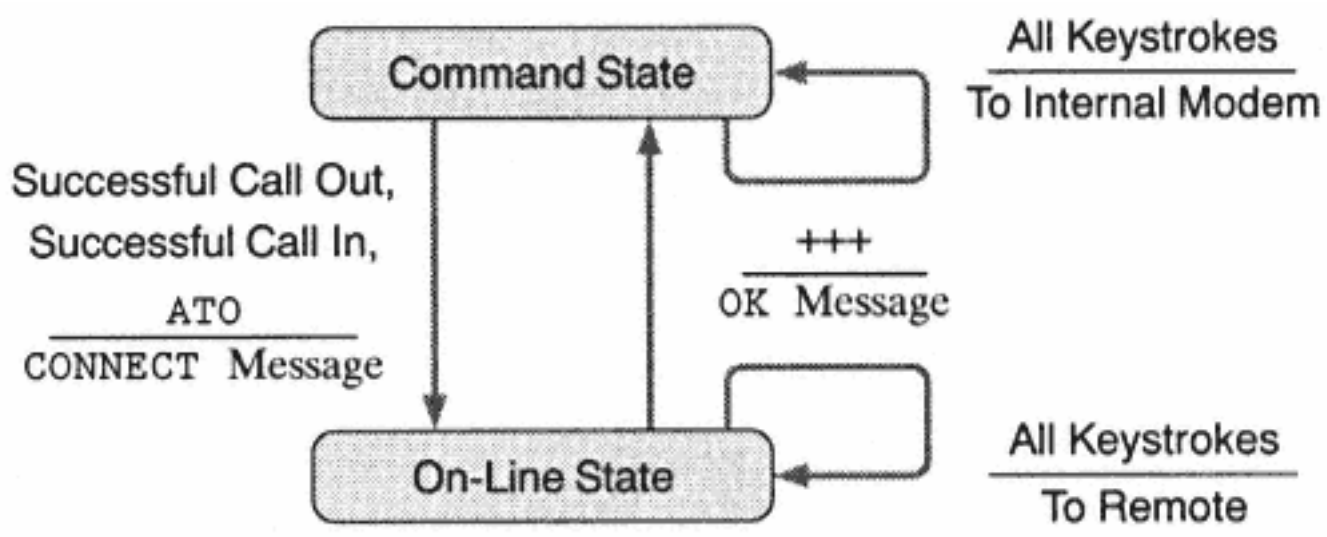

Figure 8.39
The basic state diagram for the Hayes' modem commands.

In situations where it is necessary to return control to the internal modem's command state, a sequence of three + characters must be entered (these should not be echoed; if they are, it means that the sequence has not been typed quickly enough). The response from the internal modem is `OK`, and at this point the internal modem enters its command state; keystrokes are not forwarded to the remote DTE. All characters following `AT` (or `at`) are interpreted by the internal modem; for example, to drop the carrier one types `ATH`. The sequence `ATO` returns the call to the on-line state.

The state diagram in Figure 8.39 summarizes the internal modem's actions.

Additional Hayes' modem commands and their application are given as exercises.

8.11 Summary

This chapter illustrated one way in which the distance limitations associated with standards such as RS-232-C (V.24) can be overcomenotably, by using the telephone network. Signals from a DTE cannot be placed directly on the existing telephone system without the use of a DCE (data communications equipment or data-circuitterminating equipment) commonly called a modem. Modems are responsible for modulating the signals from the DTE into a frequency that is acceptable to the telephone system, and then demodulating the telephone frequencies back into a signal acceptable to the DTE. A variety of modulation techniques exist, including amplitude modulation, frequency modulation, and phase modulation. Phase modulation permits the transmission of much higher data rates through a variety of techniques, such as phase and amplitude shifting, resulting in more bits per baud.

The connection of the DTE to the DCE is also subject to standards, one of the most popular being RS-232-C. RS-232-C has limited features, most of which are ignored by manufacturers and users. RS-232-C modem control is limited as well because there is no mechanism to permit end-to-end modem communications.

8.12 Exercises

1. (A) Show that the complete null-modem cable (shown in Section 8.4.7) exhibits the same functions as a half-duplex modem.

2. (A) Complete the constellation diagram of the ITU-T Recommendation V.29 at 9600 bps shown in Figure 8.27.

3. (A) Find a PC with an external modem. What RS-232-C connections does the PC support? What RS-232-C connections does the modem support? Can the PC control all of the modem connections?

4. (B) Obtain two PCs, two modems, and two telephones. Try running `external.c`. The software was developed for a Bell-212 modem, so it might not work with the modem(s). If it doesn't, check the modem's instruction manual and modify `external.c` so that the software can control the modem properly.

5. (C) As soon as `do_modem()` enters the `DISCONNECTING` state, all incoming characters are lost because only messages from `CLKIH` are processed. Extend `do_modem()` so that any keyboard characters that arrive while `do_modem()` is in the `DISCONNECTING` state are displayed. It may be necessary to lengthen the time between entering the `DISCONNECTING` state and issuing the `CLEAR_CALL` to allow testing of the software.

6. (C) Rewrite `external.c` so that when the call is cleared, the software waits for another call instead of terminating.

7. (B) Obtain a PC with an internal modem and the telephone number of a remote DCE/DTE. Then, running `internal.c`, call the remote DCE/DTE using the sequence shown in the text. If the telephone only supports pulse dial (as opposed to tone), prefix the number with `ATDP` rather than `ATDT`. Experiment with the commands:

`+++`. Enter command state.
`ATO`. Return to on-line state.
`ATH`. Hang up the telephone.
`ATLx`. Change speaker volume (`x=0:off`; `x=3`: loudest).
`ATEx`. Turn on (`x = 1`) command echoing; turn off (`x = 0`).

Remember, `AT` commands can only be issued while in command state. What happens if `ATO` is issued without a call in progress? It may be necessary to wait up to about 30 seconds for a response.

There are literally dozens of `AT` commands; check the user's manual supplied with the modem for more commands with which to experiment.

8. (B) Obtain two PCs with internal modems attached to the telephone network. Put one of them into auto-answer mode (i.e. it will answer the call automatically). Auto-answer is set up by changing the `S0` register to a value greater than `0` (the `S0` register selects the number of rings after which the modem will answer; `0` means disable auto-answer). For example, to auto-answer after three rings:

```
ATS0=3
```

(the command must be followed by ENTER). The value of any S register can be obtained by typing ATSx?, where x is the register number. It may be necessary to enable DTR before the call will be accepted; this can be achieved by typing

```
AT&D0
```

which instructs the modem to ignore the DTR signal.

Now, from the other PC dial the PC that is in auto-answer mode. Describe what happens. At which point can data transfer begin?

9. (C) Modify `internal.c` so that a telephone number can be entered on the command line, rather than the user having to explicitly enter it. For example, to dial 453-8100 at 2400 bps, one types:

```
C:\> internal 2400 453-8100
```

10. (C) Add modem software to the file transfer software; try transferring files to remote PCs.

11. (C) Is it possible to write software that emulates a break-out box? That is, given the signals that the serial port supports, can we monitor and display the signals and at the same time forward them (perhaps out the "other" portsomewhat like the line analyzer)? The short answer to this problem is, yes it can be done; however, the model fails in certain conditions. Write software to emulate the break-out box and identify the conditions in which the software cannot operate as a "true" break-out box.

This question requires modification of the serial port modem status change software in `commkit.c`. *Hint*: consider the issues surrounding the forwarding of modem status signals. Remember to take a copy of `commkit.c` before making any changes to it.

9
Multiplexing

9.1
Introduction

When people work at a terminal they often pause, look around, think, scratch their head, turn the page of the material they are typing, or whatever: the end result is that the channel connecting the terminal to the central computer is idle for a considerable amount of time. Even data entry clerks who can type hundreds of words a minute cannot keep the channel busy for any more than a fraction of its total capacity. As an example, the number of 10-character words a typist would have to type each minute to keep a 9600-bits-per-second line full is

$$\frac{9600\ bits\ per\ second \times 60\ seconds\ per\ minute}{10\ bits\ per\ byte \times 10\ bytes\ per\ word} = 5760\ words\ per\ minute$$

From this simple example, it is clear that a channel devoted to a single terminal is typically idle much more than it is active. However, if the channel could be *shared* among several terminals, the number of individual cables from users' DTEs to the central computer could be reduced dramatically, cutting down on the expense of cable laying and maintenance. Sharing of a single communication channel among several DTEs is known as multiplexing.

In addition to reducing cabling costs, multiplexing can also reduce the costs associated with telephone access of remote DTEs. For example, consider the costs to a company that has five telephone lines dedicated to data traffic; it may be possible to cut costs by up to 80 percent if the five lines can be replaced by a single, multiplexed line.

9.2
Background

Multiprocessing DTEs such as timesharing computers are designed to support literally dozens of remote DTEs (PCs, terminals, and so on). Each remote

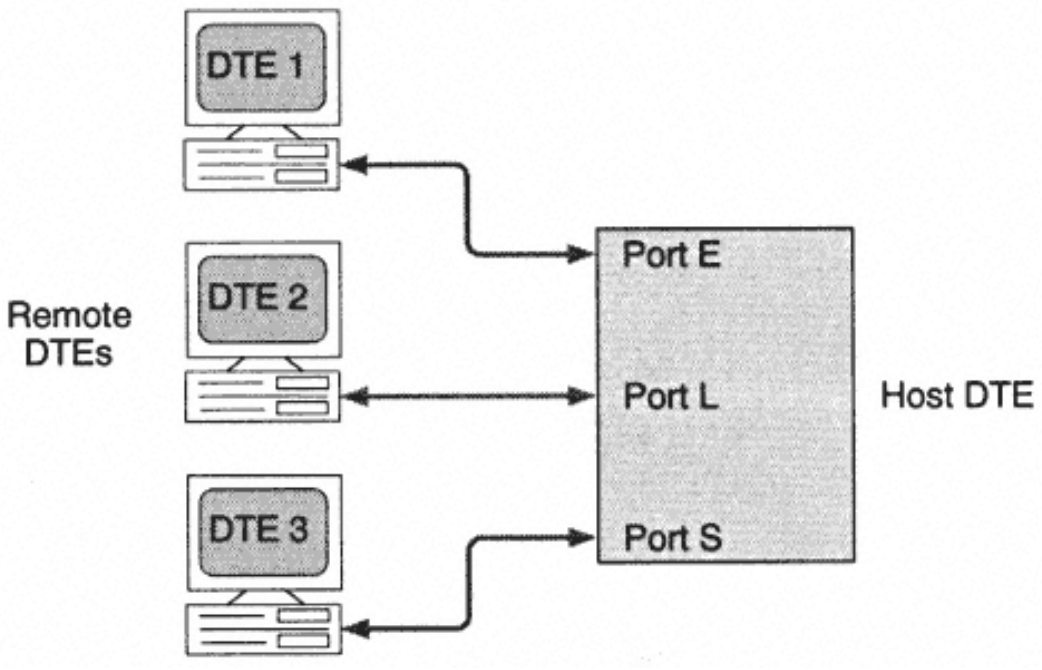

Figure 9.1
Remote DTEs connected to dedicated ports.

DTE connection is made through a port, an interface to the central computer. Most operating systems allow the various ports to be configured by the system manager or the operating system to specific terminal types (for example, dialin, asynchronous, synchronous).

A typical set of DTE connections in a multiprocessing environment might consist of the configuration shown in Figure 9.1, where three interactive terminals (DTE 1, DTE 2, and DTE 3) are connected by a series of individual asynchronous channels to ports E, L, and S, respectively, on a shared central processor (the host DTE). If the individual channels connecting the interactive terminals to the ports on the host are only partly in use, it may be possible to have some or all of the terminals multiplex their communications onto a single channel, as shown in Figure 9.2.

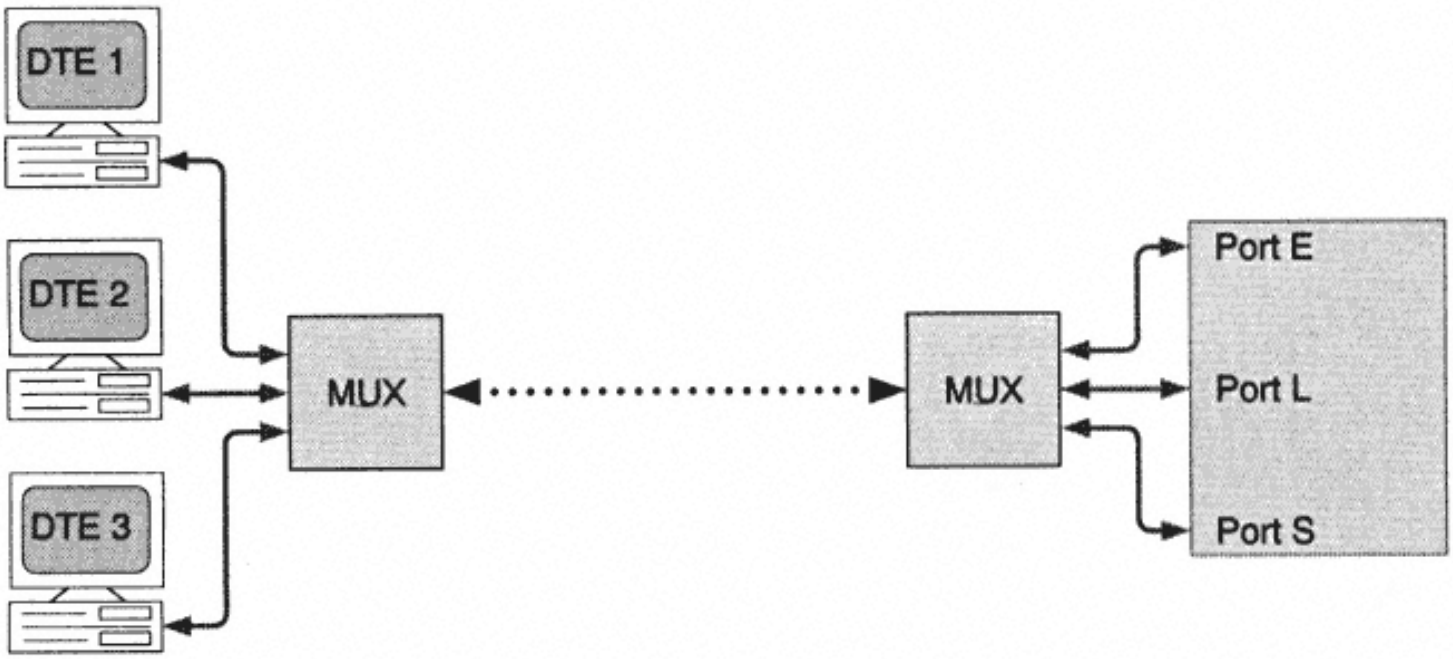

Figure 9.2
Remote DTEs sharing a channel by multiplexing.

The equipment controlling the multiplexing is known as a multiplexer (or MUX) and performs two operations:

It takes information from the various DTEs and puts it on the multiplexed channel. This is known as multiplexing.

It takes information from the multiplexed channel and supplies it to the intended destination DTE. This is known as demultiplexing.

The operation of the multiplexer should be transparent to the DTEs being multiplexed. In this situation, transparency refers to the effect of the multiplexer on communications. For example, there should be no detectable difference in communication speeds whether or not the multiplexer is present. Similarly, the data sent should not be altered in any way by the multiplexer.

The multiplexed channel is typically full-duplex, thereby allowing some DTEs to transmit information while others receive. However, the connections between a multiplexer and its multiplexed DTEs can be full-duplex, half-duplex, simplex, or all three. It is also possible for the multiplexer to support a mixture of synchronous and asynchronous DTEs. In some cases, the connection between the DTEs and the multiplexer is asynchronous and the information sent between multiplexers is synchronous. In short, the characteristics of the multiplexed channel should not have any bearing on the information that is being multiplexed. As with all other communicating devices, for the communications to succeed between a pair of multiplexers, both must agree to a common protocol.

9.2.1
Multiplexer Internals

A multiplexer is another example of the input-processing-output cycle. The multiplexing task consists of

1. Checking a connection for data.

2. Formatting the data according to the protocol used by the channel.

3. Transmitting the data on the channel.

The demultiplexing algorithm is somewhat similar:

1. Reading the data from the channel.

2. Determining the destination of the data.

3. Forwarding the data to the specific connection.

Physically, the multiplexer consists of two parts: the connections to external devices and a connection to the remote multiplexer. The connections used depend upon the distance from the multiplexer to the external devices and the remote multiplexer, meaning that multiplexers usually support both direct connections and modems.

9.3 Multiplexing Techniques

The multiplexer is responsible for ensuring that the information supplied by one DTE arrives at the correct destination DTE. It employs one of three techniques: frequency division multiplexing, time division multiplexing, and statistical multiplexing.

9.3.1 Frequency Division Multiplexing

Frequency division multiplexing, or FDM, involves dividing the multiplexed channel into a number of unique frequencies, each one assigned to a pair of communicating entities. FDM can be achieved only if the available bandwidth on the multiplexed channel exceeds the bandwidth needs of all the communicating entities.

Whenever a multiplexer receives data for transmission, the data is transmitted by it on the frequency allocated to the transmitting entity. The receiving multiplexer forwards the information received on a specific frequency to the destination associated with that frequency.

The following example illustrates how a frequency division multiplexer connects DTEs 1, 2, and 3 with ports E, L, and S, respectively, on a central host. The frequency allocation is given in Table 9.1.

The 1000-Hz separation between the channels is known as the guard band and is used to ensure that one set of signals does not interfere with another. Diagrammatically, the connections and their frequencies are shown in Figure 9.3.

The advantage of FDM is that each DTE is assigned a unique frequency that can be treated as an unshared channel. However, FDM is not widely used in data communications because of the costs of hardware that can transmit and receive signals on a variety of frequencies. An everyday example of FDM is cable television, in which many signals are "stacked up" and transmitted simultaneously over the cable. The user selects a viewing channel by tuning to that channel's frequency.

Table 9.1 Frequency allocation for Figure 9.3.

DTE-Port Pair	Frequency (Hz)
1 and E	1000014000
2 and L	50009000
3 and S	04000

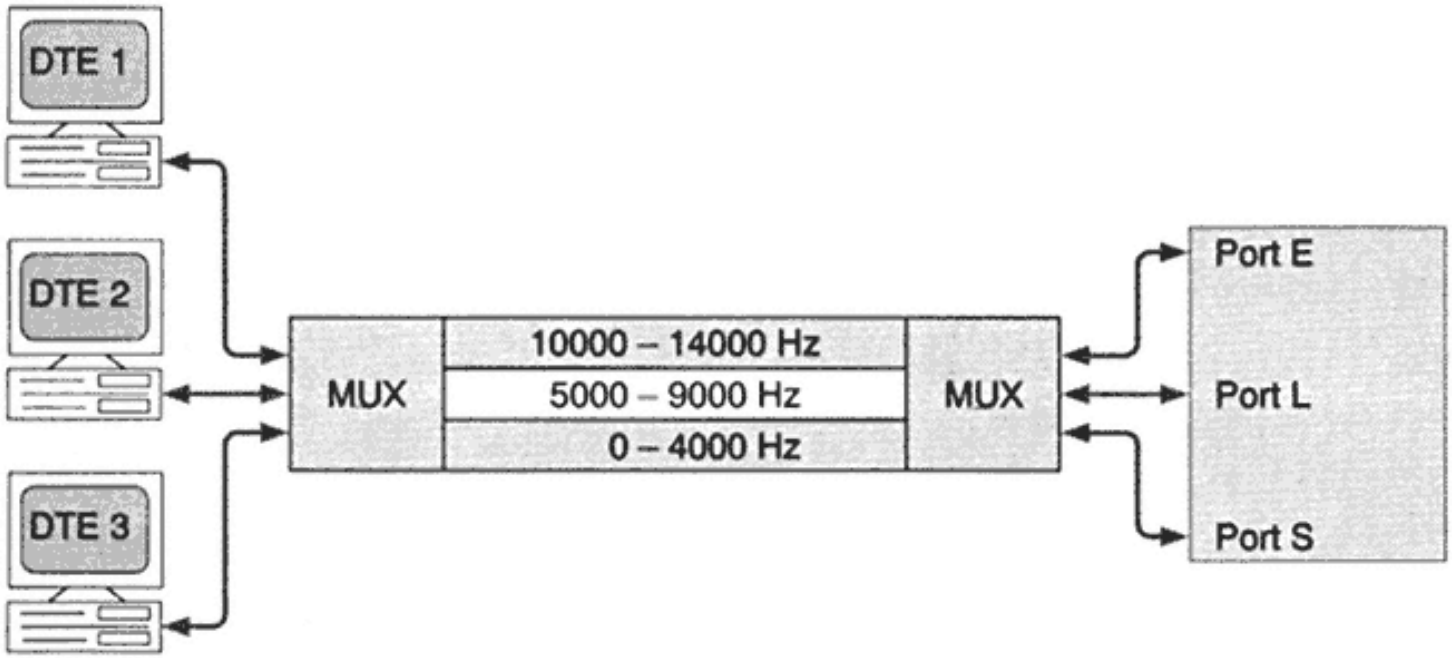

Figure 9.3
Frequency division multiplexing.

9.3.2
Time Division Multiplexing

Time division multiplexing, or TDM, requires the multiplexer to timeshare the channel between the various DTEs involved in the communication. That is, at a specific moment the remote multiplexer will send a byte from, say, DTE 1; at the next instance a byte from DTE 2 will be sent, and so on until all DTEs have been polled; the cycle is then repeated.

In Figure 9.4, DTEs 1, 2, and 3 are in communication with ports E, L, and S, respectively. The multiplexed channel is full-duplex and shared among all the communicating DTEs (the identifier in each box in the multiplexed channel between the two multiplexers should be read as "data from" the specific port or DTE; that is, it contains data rather than the identifier).

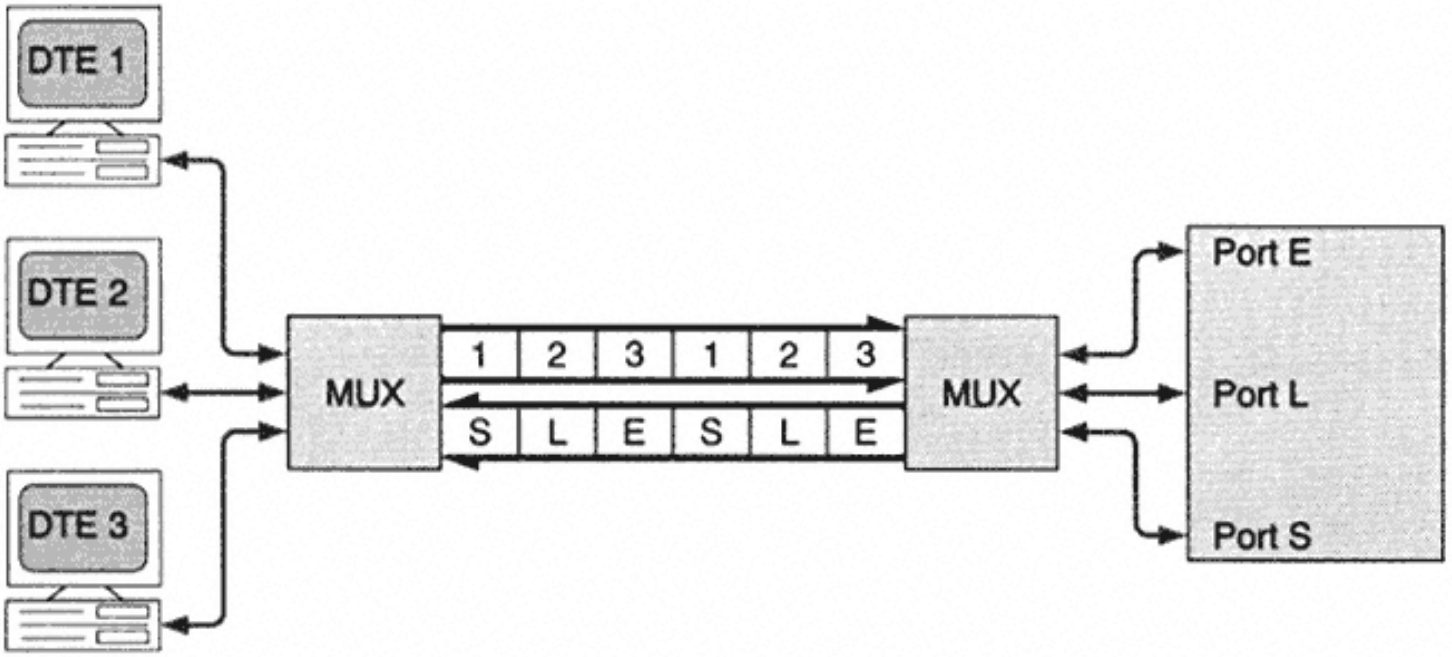

Figure 9.4
Time division multiplexing.

In Figure 9.4, the topmost channel contains data from DTEs 1, 2, and 3. The rightmost multiplexer is about to receive a byte from DTE 3, and this is followed by a byte from DTE 2. On the lower channel, data is sent from ports E, L, and S; the leftmost multiplexer is about to receive a byte from port S.

To prevent information arriving at the wrong DTE, both multiplexers must be synchronized. That is, the bytes must be sent in an agreed-upon order and each byte that is received must be for the specific DTE (or port) for which it was intended. Synchronization can be achieved in a number of ways. A common approach is to use a special bit pattern to indicate the start of a new cycle, so that if *N* DTEs are being multiplexed, the (*N* + 1)th byte to be transmitted is a special, synchronizing bit pattern.

A time division multiplexer polls each DTE to determine if there is information to be sent. However, a problem arises if one of the DTEs has nothing to send: what does the multiplexer transmit? Something *must* be sent because not sending a byte means that the time allotted to one DTE may be used by another, potentially resulting in the information arriving at the wrong destination. For example, if DTE 2 has nothing to send to port L, data from DTE 1 might be sent instead.

This problem can be overcome in several ways, including

Having a reserved bit pattern (for example, the *NUL* character) that is sent whenever a DTE (or port) has nothing available for transmission.

Transmitting *nine* bits between the multiplexers, eight for data and the ninth to signal whether the byte contains data or is empty.

9.3.3 Statistical Multiplexing

Statistical multiplexing attempts to overcome the problem of idling DTEs by sending information from a DTE only when it is available. For example, if three DTEs are sharing a channel and only one is active, then as much of the channel as possible should be given over to the active DTE. (In the time division multiplexing example in the previous section, two of the three time slots would be empty because of the inactive DTEs.)

Since the multiplexers must be able to determine the intended destination of each byte received, a unique identifier either of the source (i.e., the sending port or DTE) or of the destination (i.e., the receiving port or DTE) of the byte must be included with each byte. This has two implications:

The bandwidth is reduced because the identifier is sent with each byte. If the identifier is a byte (a reasonable choice, since the multiplexed channel is probably byte-oriented), the bandwidth is halved.

The multiplexer must know the destination of the byte. Each byte must be transmitted with an identifier that can indicate either the byte's source or its intended destination. If the identifier is the source's address, the

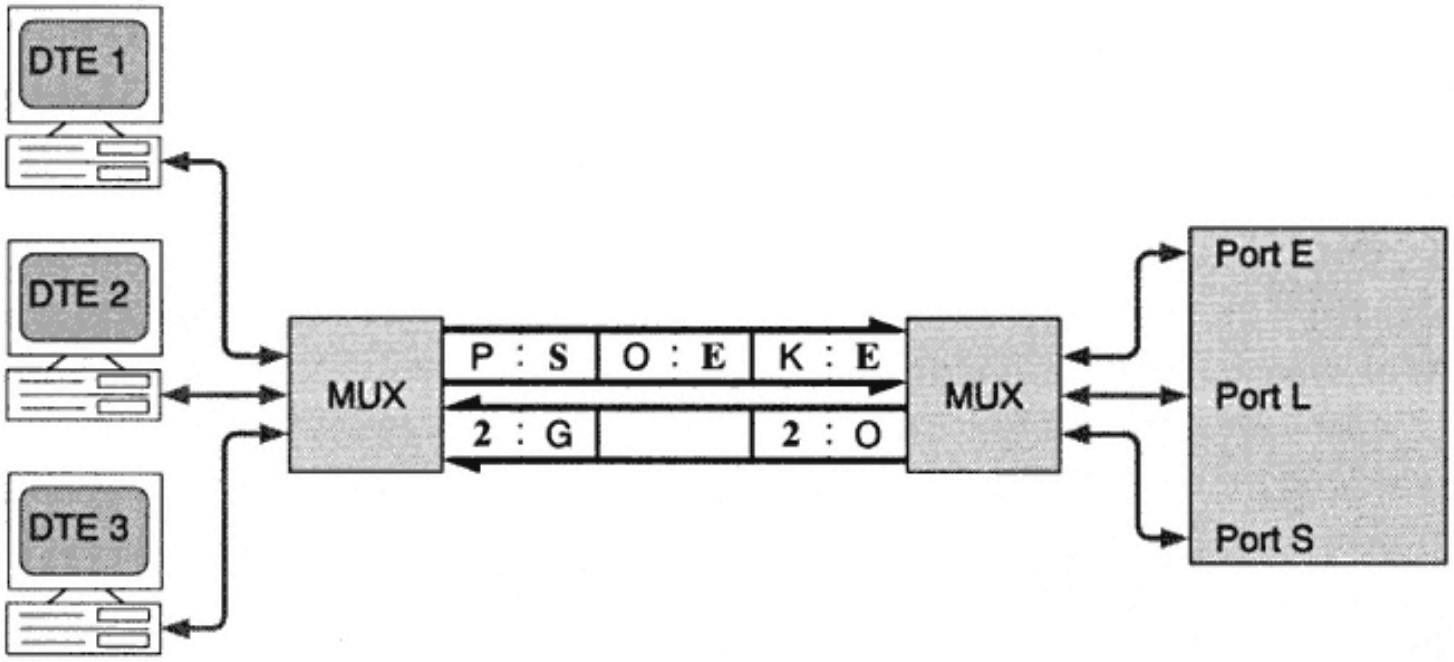

Figure 9.5
Statistical multiplexing.

remote multiplexer must map that address into a destination connection. Similarly, if the identifier is a destination address, the local multiplexer must map the local device's identifier into the destination address. Either of these approaches can be through the use of mapping tables.

In addition to the reasons already noted, statistical multiplexing differs from time division multiplexing in that if none of the DTEs are active, the multiplexed channel will be idle.

Figure 9.5 shows how a pair of statistical multiplexers can function. As bytes are sent, they are prefixed with an identifier that allows the receiving multiplexer to determine their destination. In this example, each data byte is sent as two bytes: the address of the destination (displayed in bold) and the data.

In Figure 9.5, the topmost channel (flowing from left to right) contains two bytes destined for port E (K followed by O) and one byte for port S (a P). The lower channel has two bytes for DTE 2 (G followed by O); additionally, there is a period during which nothing is being transmitted on the channel.

If all devices connected to a multiplexer transmit simultaneously for a sustained period, there may not be sufficient bandwidth to handle all of the traffic. For these situations, the multiplexers may simply discard the additional input or buffers may be used to hold the data until transmission can take place.

9.4
Port Selectors

In most multiprocessing systems, there are more users than there are possible connections to the central computer; in some cases, there may be more remote DTEs than there are ports to the computer. From the system manager's

point of view, this is quite reasonable since few users want to be on the computer 24 hours a day (there are exceptions), and the computer manufacturer may place a limit on the maximum number of ports that the computer can support.

In situations where there are more remote DTEs than ports on the central host, a device known as a port selector (or front end, switch, or terminal concentrator) is employed to manage the connections from the remote DTEs to the host's ports. The port selector supports two sets of connections: those to the remote DTEs and those to the host's ports. The port selector is typically connected to *all* possible remote DTEs and to all ports on the central host.

When a user on a remote DTE wants to initiate a communication with the central host, some form of signal is sent by the user to the port selector for example, a series of characters (typically one or more carriage returns), a break indication, or the detection of a carrier signal generated when the DTE is powered on. If a free port on the central host is available, the port selector makes a *logical* connection between it and the remote DTE. Thereafter, all communications between the remote DTE and the host's port are handled by the port selector mapping the data from the DTE to the port or vice versa. In Figure 9.6, a central host has two ports (A and B) and there are four possible remote terminals (DTE 1 through DTE 4); the port selector has mapped DTE 1 to port A and DTE 3 to port B.

When the communication finishes, the logical connection is broken and another DTE can use the port. Over a period of time, a DTE might be involved in any number of communications with the central host; however, each communication may use a different port. If all ports are in use, the port selector ignores the incoming requests (possibly issuing a diagnostic message to the user at the remote DTE).

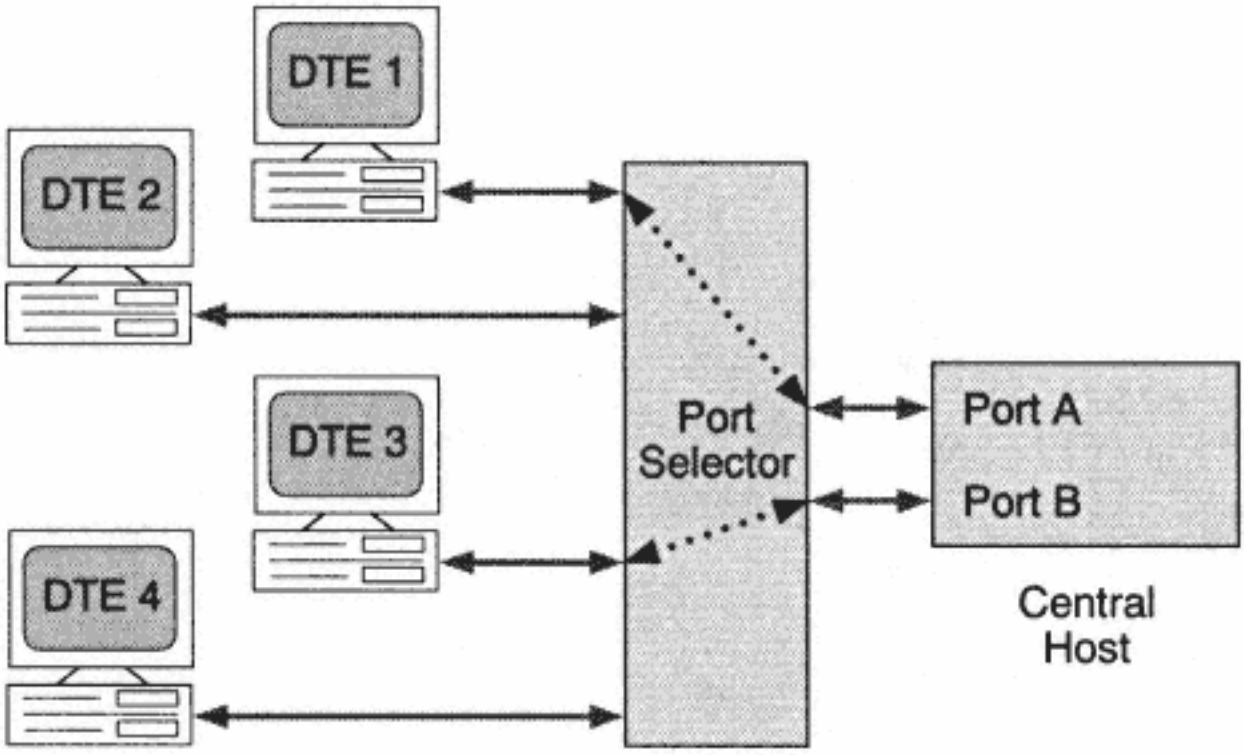

Figure 9.6
A port selector.

Port selectors can offer multiplexing capabilities as well. That is, one or more remote DTEs might share a single channel from a remote site; the port selector will demultiplex the channel and make logical connections to free ports using the same techniques as for the directly connected remote DTEs.

9.4.1
Other Multiplexer Applications

Until now the examples of multiplexing showed one multiplexer connected to a group of remote DTEs and a second multiplexer connected to a central host. In some situations, it may be necessary to have the several central hosts (rather than one). For example, in Figure 9.7 a pair of multiplexers support two central hosts and several remote DTEs.

The connections between the various DTEs can be:

Static. If the connections are static, the multiplexers "know" the destination of each byte sent. Static connections can be implemented using statically allocated mapping tables, meaning, for example, that each byte from DTE 2 is always sent to DTE S (and vice versa).

Dynamic. Static connections force users to seek out terminals that connect to specific DTEs. By making the connections dynamic using some form of a port selector, users can choose their destination DTE.

When the remote user is allowed to select one of several central DTEs, it is necessary to indicate the required destination to one of the multiplexers. Once the destination is known, the information can be used by the multiplexers to establish a path from the user's DTE to the central DTE. If dynamic connections are allowed, the multiplexers may be required to support a protocol whereby the connection information (i. e., the source and destination addresses of the devices) is exchanged.

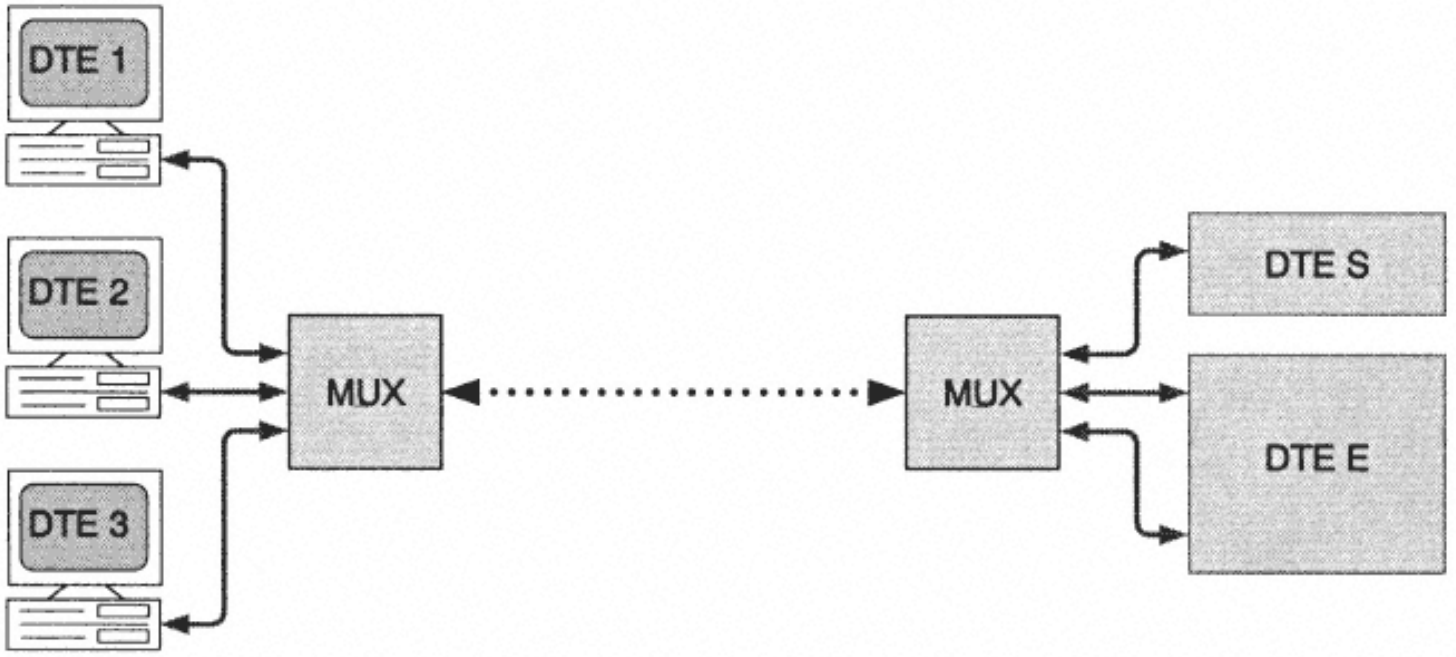

Figure 9.7
Multiplexing with multiple central hosts.

9.5 Multiplexing with Commkit

Commkit can be used to illustrate multiplexing across a single channel between two PCs. The Commkit diskette is supplied with an example of statistical multiplexing between three separate processes. Time division multiplexing and frequency division multiplexing can also be implemented using Commkit; they are left as exercises.

9.5.1 Hardware Requirements

The multiplexing software is intended to be run on two PCs interconnected by a serial port. The traffic associated with the statistical multiplexers can be monitored if a line analyzer is included; see, for example, Figure 9.8.

9.5.2 Design

The statistical multiplexer software is designed to solve the following problem:

Two PCs should be interconnected so that three pairs of processes can communicate over a single channel using statistical multiplexing. Each process will accept local messages and forward them to the remote process, and at the same time receive and display messages from the remote process. A message is a single byte.

There are three processes. The first will accept local keyboard information for transmission and accept remote keyboard information for display. The second process will transmit a byte every 18th of a second while displaying bytes from the remote process. The final process will transmit one byte once a second and display the byte sent by its remote process.

The output displayed on a local PC is the characters generated and transmitted by the remote PC. Keyboard output is displayed on lines 0 through 6, and output from the second process on lines 8 through 14; lines 16 through 22 display the output from the final process, sent once a second.

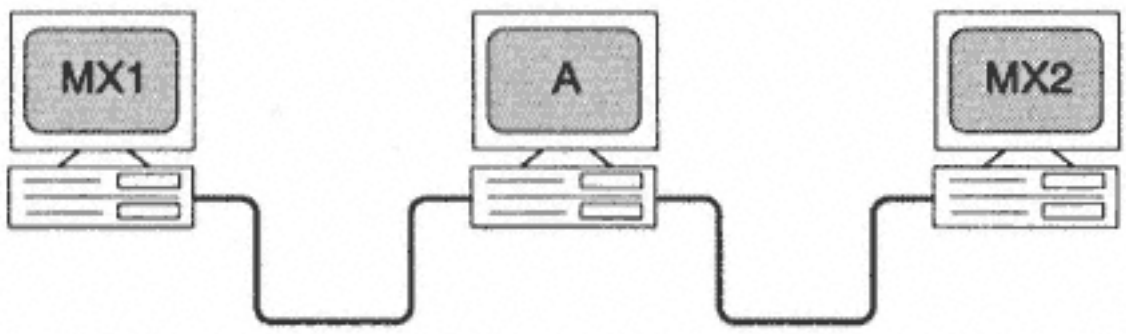

Figure 9.8
The Commkit multiplexer hardware configuration.

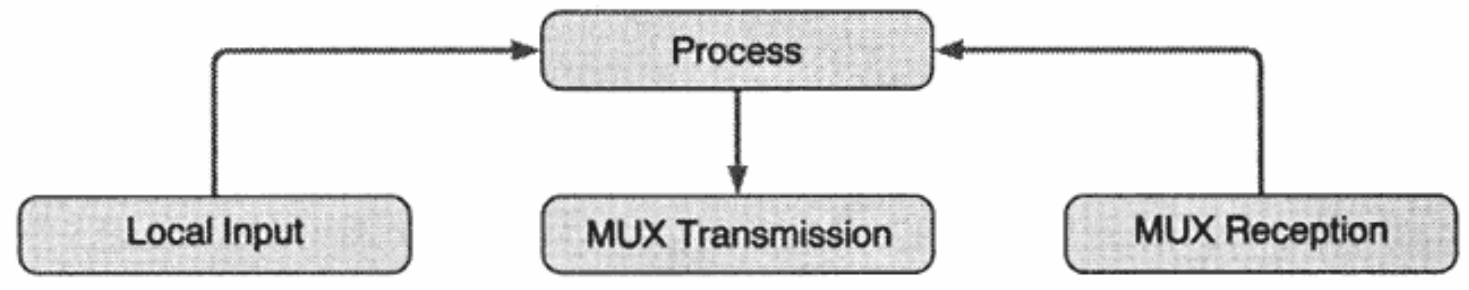

Figure 9.9
The structure of the Commkit statistical multiplexer.

Process Design

From the description of the problem, one can design a "generic" process that accepts local data and forwards it for transmission and that accepts remote data for display; such a process is shown in Figure 9.9.

Local input comes from two devices: the keyboard interrupt handler, `KEYIH`, and the clock interrupt handler, `CLKIH`. The clock interrupt handler generates two messages: one every 18th of a second, the other once a second. Each of the three processes sends a message to the local MUX for transmission to the remote MUX. At the same time, messages that are received by the local MUX will be forwarded to the destination process for display.

The Transmission MUX

The transmission multiplexer is responsible for forwarding two bytes: the first is an identifier signifying the intended destination process, and the second is the data. The multiplexer can be represented as the state machine shown in Figure 9.10.

The MUX transmission process will stay in state `IDLE` until a byte is available for transmission, at which point the identifier of the intended destination (`DST`) is transmitted and the state changes to `DST_ID`. As soon as

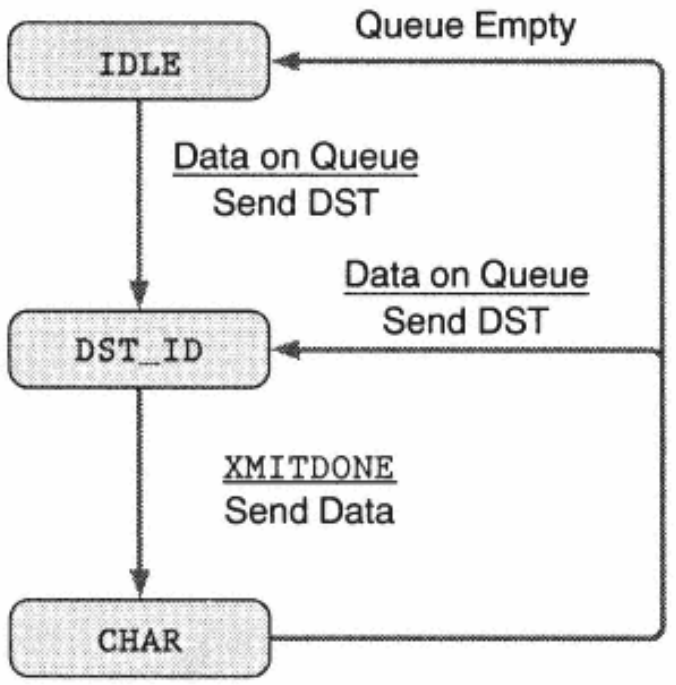

Figure 9.10
The Commkit Multiplexer Transmission State Machine.

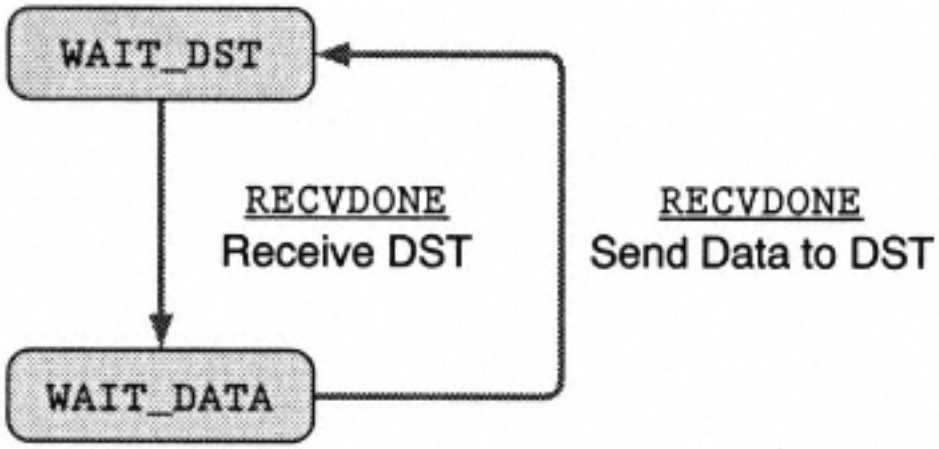

Figure 9.11
The Commkit Multiplexer Receive State Machine.

the transmission of the identifier has completed, control is passed to state `CHAR` and the data is sent. When this transmission has completed, the queue is checked again. If there is data on the queue, the destination identifier is sent and control is passed to state `DST_ID`. If there is no data on the queue, the state returns to `IDLE`.

The Reception MUX

The reception multiplexer must wait for two bytes: the first contains the destination process identifier, and the second carries the data to be displayed by the specified process. In Figure 9.11, a state machine consisting of two states represents the reception multiplexer.

The reception MUX enters the `WAIT_DST` state and waits for input. The first byte received is assumed to identify the destination process (`DST`) and the state changes to `WAIT_DATA`. The second byte received is the data (`Data`). When it is received, it is forwarded to the process indicated by the first byte received. Control returns to the `WAIT_DST` state.

9.5.3
Implementation

The implementation of the Commkit statistical multiplexer is based upon the design described in the previous section. Implementation requires the use of three processes as well as all of the low-level interrupt handlers. The processes and their associated queues are as follows:

`do_mux (). Do_mux()` is the foreground process, accepting single bytes sent to the `APPLICATION` queue from the `KEYIH` interrupt handler. These bytes are sent to the serial port responsible for emulating the multiplexer (indicated by `handler_id`).

Bytes received from the channel by the multiplexer process for `do_mux()` are put on queue `APPLICATION` until they are received for display.

`do_18_clk()`. This process receives messages 18 times a second from `CLKIH`. After receiving a message it sends the "next" character from the string `abcdefghijklmnopqrstuvwxyz` to the queue associated with

the multiplexer serial port (handler_id). A single character is sent each time a message from CLKIH is received, starting at a and proceeding one character at a time to z. When the end of the string is reached, the cycle is repeated.

The do_18_clk() process is associated with queue BACKGROUND_1. This process displays all bytes received from the channel and forwarded by the local multiplexer to BACKGROUND_1. Do_18_clk() displays the bytes received from this queue.

do_l_clk(). Do_1_clk() is associated with queue BACKGROUND_2, receiving messages once a second from the clock interrupt handler, CLKIH. Upon receipt of a message, it sends the "next" byte from the string 1234567890; when the end of string is reached, the cycle is repeated. Messages sent to the BACKGROUND_2 queue by the multiplexer are displayed by do_1_clk() after they are received.

Data is sent to and received from the remote multiplexer through one of the serial ports (specified by the user). Low_level() is called whenever a message is sent or when a byte is received:

Transmission MUX. The transmission MUX is an implementation of the transmission multiplexer described in the Process Design section. It is implemented in low_level() and mux_send(). There are two conditions under which a transmission can occur:

1. If the multiplexer is idle (indicated by mux_send_state having a value of IDLE) and low_level() is called with a code of MSG_AVAIL. Mux_send() is called at this point and the transmission commences. If the multiplexer is not idle, the message remains queued.

2. If the multiplexer has completed a transmission of a data byte (indicated by a code of XMITDONE), the multiplexer state is CHAR, and if data is on the multiplexer's queue, a transmission will occur.

Reception MUX. The reception multiplexer, mux_recv() is an implementation of the reception multiplexer state machine described in the Process Design section. Once the destination identifier and the byte have been received, the byte is forwarded to the destination process; the flow is shown in Figure 9.12.

To allow the user the opportunity to distinguish between messages sent by the various processes, the PC's screen is divided into thirds: the top third (lines 0 through 6) displays the data received from the remote keyboard process; the middle third (lines 8 through 14) shows the data received from the remote 18-times-a-second process; and the lower third (lines 16 through 22) displays the data received from the remote once-a-second clock process. When a process's screen is filled with information, it is cleared and output resumes in its top left corner.

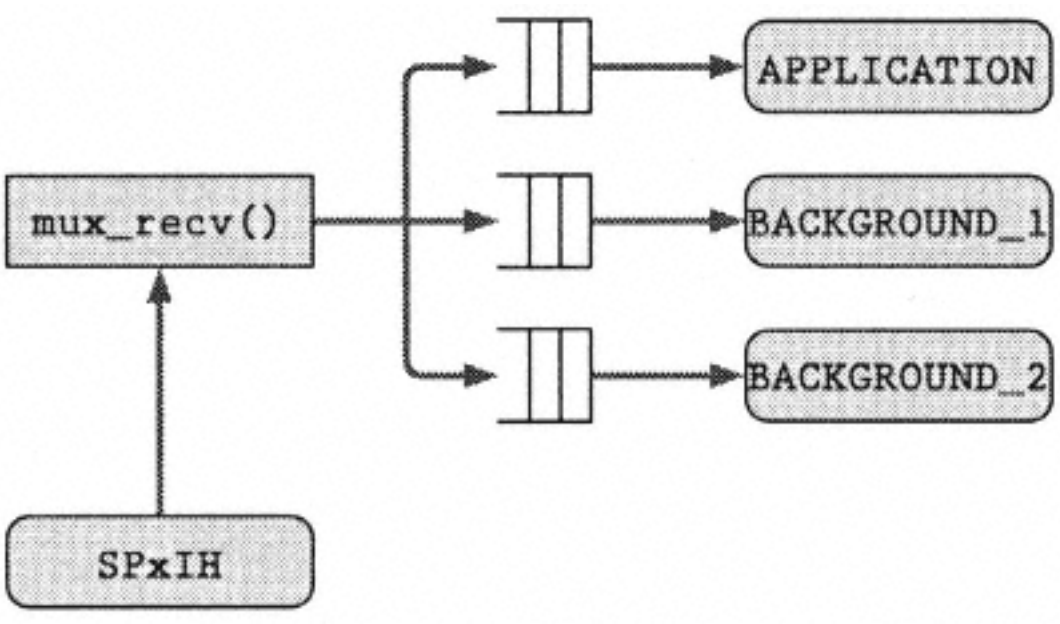

Figure 9.12
Multiplexer reception information flow.

The identifier associated with the intended destination is simply the identifier of the transmitting process. For example, keyboard characters sent from `BACKGROUND_1` on one PC are received by the `BACKGROUND_1` process on the remote PC.

Finally, there is a problem in the implementation as it now stands: quite simply, when both multiplexers are started, one may begin transmitting before the other. This may well cause a loss of data; more important, however, neither MUX can tell if the byte received is a data byte or a destination identifier. This means that the two multiplexers may be out of synchronization, with one (or possibly both) treating the data as the destination identifier and the destination identifier as the data.

The solution adopted in the multiplexer software has been to set the eighth bit when transmitting the identifier and to clear it when sending data. Then, when a byte arrives at a MUX it is possible to immediately determine whether the byte is a destination identifier or a data byte simply by checking the eighth bit.

9.5.4
Compiling and Using the Statistical Multiplexer

The statistical multiplexer is written in Turbo C and can be found in `statmux.c` on the Commkit diskette; an executable version, also on the diskette, can be created using the `make` utility:

```
C:\> make statmux.exe
```

`Statmux.c` includes `general.h`, `ascii.h`, and `devices.h`. Once compiled, `statmux.obj` is linked with `commkit.obj` and `srcalls.obj` to create the executable `statmux.exe`. The executable version can then be run (a compiled version of `statmux` is supplied on the Commkit diskette);

Table 9.2 Destination bytes as shown by the line analyzer.

Destination	Byte Value	Displayed as
APPLICATION	0x84	à
BACKGROUND_1	0x85	å
BACKGROUND_2	0x86	ç

the line speed (50, 300, 1200, 2400, 4800, or 9600) and port number (1 or 2) must be specified. For example, to start the statistical multiplexer at 1200 bps using port 1, type:

```
C:\> statmux 1200 1
```

Once the program is running, the screen is cleared and divided into thirds by a pair of dashed lines. When both PCs are connected and the statistical multiplexers are running, anything generated by a process on one PC is displayed on the other in its corresponding part of the screen.

If the line analyzer is used, pairs of bytes will appear on the analyzer screen. The first byte is the destination process, and the second byte is the data. The destination bytes are listed in Table 9.2.

If the line speed chosen is too slow, the message queues may be dumped since data is being supplied to the serial port faster than it can be sent. Communication will continue but in a very degraded fashion. The problem can be remedied by selecting a higher line speedwhich line speed is left as an exercise.

9.6 Summary

This chapter examined another method of reducing line charges, notably the sharing of a channel among several DTEs or processes through the use of a multiplexer. The basic function of a multiplexer is to accept data from a number of devices and transmit it on a single, shared channel while simultaneously receiving data from the channel and routing the information to the correct destination device.

There are a number of different multiplexing techniques:

Frequency division. Frequency division multiplexing requires each pair of communicating devices to be given their own frequency on the multiplexed channel. The multiplexers put data on the channel using the frequency associated with the devices. The frequency at which data is removed indicates the destination of the data.

Time division. In time division multiplexing, data from each device is sent at a precise time interval on the multiplexed channel. If the device has nothing to send, a signal is sent to this effect. The receiving multiplexer is synchronized with the transmitter to ensure that the received data is supplied to the correct destination device.

Statistical. Statistical multiplexing attempts to reduce the amount of wasted bandwidth associated with time division multiplexing by transmitting information only when a device has it to send. To allow the receiving multiplexer to determine the intended destination of the data, an identifier is sent with every data byte.

9.7 Exercises

1. (A) Describe how device polling works. Compare device polling with multiplexing.

2. (A) Two approaches to handling idle connections were described for time division multiplexing. Describe the advantages and disadvantages of each.

3. (AB) Determine both analytically and experimentally the minimum line speed needed to allow the data to pass without loss (or without crashing Commkit because of lack of queue space). The important numbers to consider are shown in Table 9.3.

If the channel is set to 9600 bits per second, approximately what percent of the channel is being used? Assume 10 bits per byte (8 bits of data, 1 start bit, and 1 stop bit).

4. (C) The statistical multiplexing example has "well-known" destinations already assigned in the software. Devise and implement an algorithm that allows the destination process to be assigned dynamically.

5. (C) Develop and implement an algorithm to support frequency division multiplexing using Commkit.

One possible solution is to treat each 8-bit byte that the multiplexer transmits as a collection of eight frequencies, each one assigned to a separate process. If there is nothing to be sent, the multiplexer sends `0xFF` (each bit position having a value of 1). If a process supplies a byte to be transmitted, the multiplexer should change the bit position associated with the process to 0 (to indicate the start of data). Then, for the remaining eight

Table 9.3 Data for Exercise 3.

Process Name	Bytes/Second
`APPLICATION`	18
`BACKGROUND_1`	18
`BACKGROUND_2`	1

data bits, the bit associated with the transmitting process contains the "next" data bit in the data to be sent. Once the data has been sent, the bit position should return to sending 1s. The reception multiplexer must scan each received byte for a non-zero bit value (the start-of-data signal). The data should be assembled from the next eight incoming bytes.

This solution requires a considerable amount of coding, perhaps analogous to the extra work required to support multiple frequencies.

6. (C) Design and implement a time division multiplexing program using Commkit to support the transmission of 7-bit data. The eighth bit can indicate whether the transmitted byte is data or is empty.

7. (C) In light of the last question, consider how to transmit 8-bit data using time division multiplexing and Commkit. It *might* be possible to use mark and space parity.

10
Wide Area Networks

10.1
Introduction

The material covered thus far in Part 3 has illustrated techniques for minimizing transmission overheads, overcoming distance limitations, and maximizing channel utilization. However, even these techniques cannot resolve the following problems:

A branch office uses a multiplexer and modem to connect to its head office located several hundred kilometers away. In order to allow instant access to the head office, a telephone connection is maintained throughout the working day (including coffee breaks, lunch time, and meetings). Although the occasional file transfer and some busy periods during the day may, from time to time, keep the channel at near capacity, there is insufficient network demand to warrant the telephone charges paid by the company.

A telephone company analyzes its network traffic and discovers that the average duration of a voice call is about five minutes while data calls can last for hours. Meanwhile, the annual growth in data traffic exceeds that of voice by roughly a 4-to-1 margin.

These two examples highlight an interesting contradiction: although more data traffic may suggest additional revenue for the telephone company, it can also result in more customer dissatisfaction because of failed call attempts (by both data and voice customers) due to congestion in the telephone network caused by the data traffic.

There are a number of possible solutions to this problem, including

Installing more telephone circuits, thereby reducing the probability of failed calls.

Offering inducements to existing and potential data customers to get them off the voice network.

The first solution, increasing the number of circuits, will alleviate the situation only for the short term, since data traffic will continue to grow. This leaves the second solution, getting data customers off the voice network, as the only real alternative. The problem is, where do they go?

10.2
Alternatives to the Public Telephone Network

10.2.1
Private Lines

A common approach to getting data customers off the voice network is for the telephone company to lease them a private line, also referred to as a leased or dedicated line. This line connects the data customer's DTE directly (via the telephone company's circuits) to the remote DTE. There are a number of advantages in having a private line, including minimal line noise (private lines are usually conditioned), guaranteed access to the telephone network, and the "freeing up" of a voice circuit. (However, with the advent of quality digital communications, the demand for conditioned lines is decreasing.)

This solution has been taken even further: some larger companies rent entire private telephone networks from telephone companies or other organizations, thereby ensuring that all their calls (both data and voice) can be placed. It is increasingly commonplace to find private companies (other than the telephone companies) with their own T1 networks (see Chapter 8).

10.2.2
Public Data Networks

In some countries, public data networks (or PDNs) have been installed, designed specifically for the transmission of data rather than voice. Circuit-switched data networks provide digital transmission facilities (i.e., the modulation and demodulation of signals is not required) for high-speed, highvolume communications. A typical digital data network may allow full-duplex communications for both asynchronous and synchronous DTEs at speeds of up to 56 kilobits per second. Many telephone companies support data networks designed specifically for digital (as opposed to voice) communications. For example, in the United States, AT&T offers Data-Phone Digital Services (DDS), while in Canada Telecom Canada offers a service known as Dataroute.

The voice network can be avoided by attaching DTEs directly to the digital network using a DCE designed specifically for digital technology.

Accessing Public Data Networks

Chapters 4 and 8 dealt with RS-232-C (V.24) and its proposed successors, RS-422-A, RS-423-A, and RS-449. To overcome the well-known limitations

of RS-232-C, the new standards increase the number of connectors from a single 25-pin D connector to a pair of them for RS-449: a 37-pin primary channel connector and an optional 9-pin secondary channel connector. For all the benefits associated with these new standards, the number of level converters (see Chapter 3) and connector pins increases the cost. What is needed is a mechanism that supports *fewer* connections but offers at least the same number of functions.

ITU-T Recommendations X.20 (asynchronous) and X.21 (synchronous) are examples of standards defining digital access to public data networks. In these standards, functions are coded as digital messages and passed between the DTE and DCE as streams of character strings using 15-pin connectors rather than a single pin devoted to a specific function. X.20 and X.21 provide services other than data transferfor example, automatic calling features and signals that allow the DTE to determine the state of a call. The X.21 circuits are shown in Figure 10.1.

In X.21, both the DTE and the DCE are defined as state machines, with the various states assigned to one of four phases: quiescent, call control, data transfer, and clearing (call) circuits. State transition is caused by changing the value of the signals on the T (Transmit), C (Control), R (Receive), and I (Indicate) circuits (T and C are associated with the DTE, while R and I are associated with the DCE).

For example, if both the DTE and the DCE are in the quiescent phase (i.e., a call can be signaled by either the DTE, as an outgoing call, or the DCE, as an incoming call), the T and R circuits idle with a continuous stream of 1s while the C and I circuits are in an off state (also signaled by a continuous stream of 1s). The DTE can signal a call request by changing the value of the T and C circuits to a stream of 0s and entering the on state (a stream of 0s), respectively. The DCE responds with either a stream of +'s indicating that the DTE can place the call, or a stream of *BEL* (ASCII `0x07`) characters, signaling that the outgoing call has collided with an incoming call. Either response is prefixed by a pair of *SYN* bytes and is sent on the R circuit (the I circuit remains in the off state).

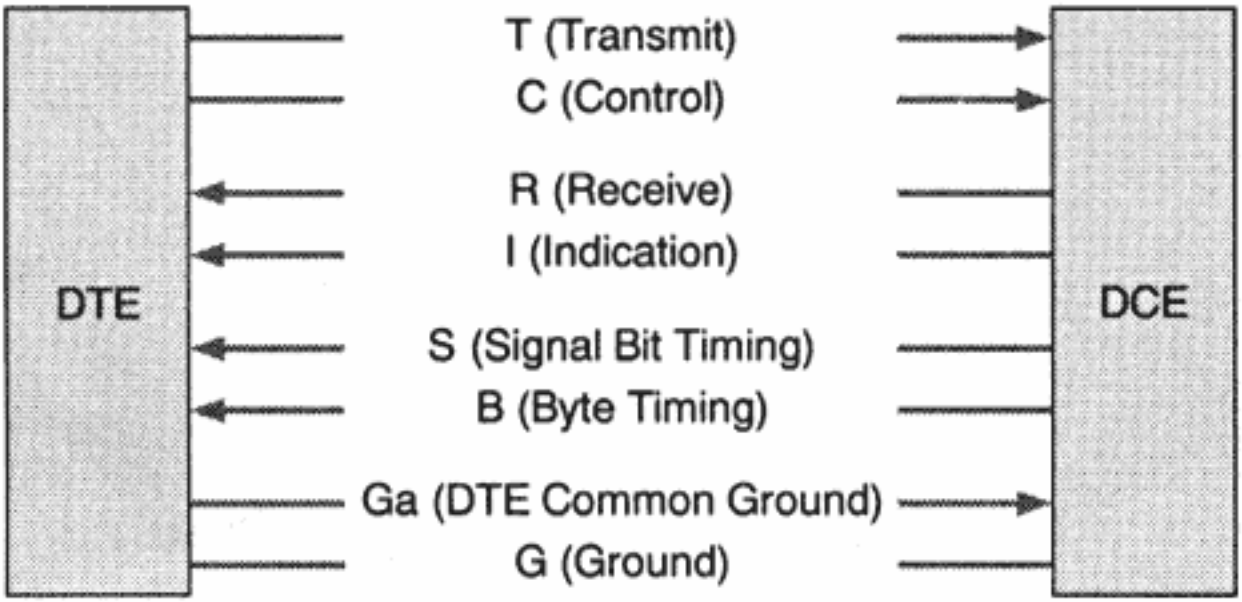

Figure 10.1
X.21 circuit definitions.

If the call can be placed, the DTE supplies the DCE with the number of the party called as a stream of ASCII (IA5) digits on the T circuit. The DTE can request specific facilities (such as making the call collect) as part of the data stream; these facilities are defined in ITU-T X.2. The final digit is followed by a continuous stream of 1s on the T circuit, after which the DTE waits for the DCE to signal the progress of the call. While the DCE is attempting to place the call, the R circuit continues to idle with a stream of + characters. The DCE can inform the DTE of the progress of the call (for example, Number Busy or Invalid Facility Request) by sending Call Progress Signals on the R circuit in place of the + characters. Call Progress Signals are defined in ITU-T X.96.

A DTE is informed of an incoming call by receiving a *BEL* on the R circuit (prefixed by a pair of *SYN* characters) while in the quiescent phase. It accepts the call by changing the C circuit from a stream of 0s to a stream of 1s.

Data transfer can commence when the T and R circuits indicate a stream of 1s and the C and I circuits are in the on state (a continuous stream of 0s). X.21 supports full-duplex communications; data is sent on the T circuit and received on the R circuit. If data is not available, the circuit (T or R) idles with 1s, while the C (or I) circuit remains in the on state.

Either DTE can clear the call by setting the T circuit to transmit a stream of 0s and returning the C circuit to the off state. The DCE connected to the DTE clearing the call responds by setting the I circuit to off. The remote DCE signals the call-clearing request to the remote DTE by setting the R and I circuits to a stream of 0s and the off state, respectively.

Both balanced (X.27) and unbalanced (X.26) circuits are defined for X.20 and X.21. The X.27 standard is similar to the RS-422-A standard and is recommended for DTEDCE line speeds greater than 9600 bps. For lower speeds, unbalanced circuits may be specified using X.26, which is similar to RS-423-A.

Two ITU-T standards, X.20*bis* and X.21*bis*. are available for DTEs that are configured to support either RS-232-C or V.24. They are designed to convert X.21 (and X.20) signals into an equivalent RS-232-C or V.24 value to allow a DTE with analog equipment access to a digital public data network.

10.2.3
Sharing Channels

The impact of data traffic on the voice network is clearly minimized through the use of private lines or public data networks. However, the problem of channel utilization still exists: there will be periods throughout the day when the channel is being paid for but is not in use.

What is needed is a communication channel that is always available to the customer, but for which the customer is charged only when transmitting or receiving information. If this channel could be shared among a number

Figure 10.2
A shared channel between
two cities, A and B.

of different customers, costs could be shared as well (based, perhaps, on the volume of data transmitted each month). For example, a number of customers in city A could share a channel to city B. Multiplexing could be used to ensure that the information sent by a customer in city B arrives at the correct DTE in city A and vice versa. Through statistical multiplexing, the channel would be used only by customers who actually required the channel at any particular moment, as depicted in Figure 10.2.

In Chapter 9, all communications were assumed to take place between one or more central hosts and a group of remote DTEs multiplexing the same channel owned by a single organization. However, in this situation there can be a number of DTEs at either end of the multiplexed channel belonging to different customers, as depicted in Figure 10.3.

Essentially, two approaches can be used to support the transfer of information between the DTEs:

All DTEs can be statically paired using a predefined connection (for example, DTE L in city A is always connected to DTE S in city B). The static connection is simple to implement and requires very little processing on the part of the multiplexers; however, users are given little flexibility.

The various DTEs dynamically select their intended destination (for example, the user of DTE I in city A may want to connect to DTE P, Q, or R). The dynamic connection is more difficult to implement for a variety of reasons: mechanisms are needed to detect whether the remote DTE is available as well as to determine if both DTEs are owned by the same cus-

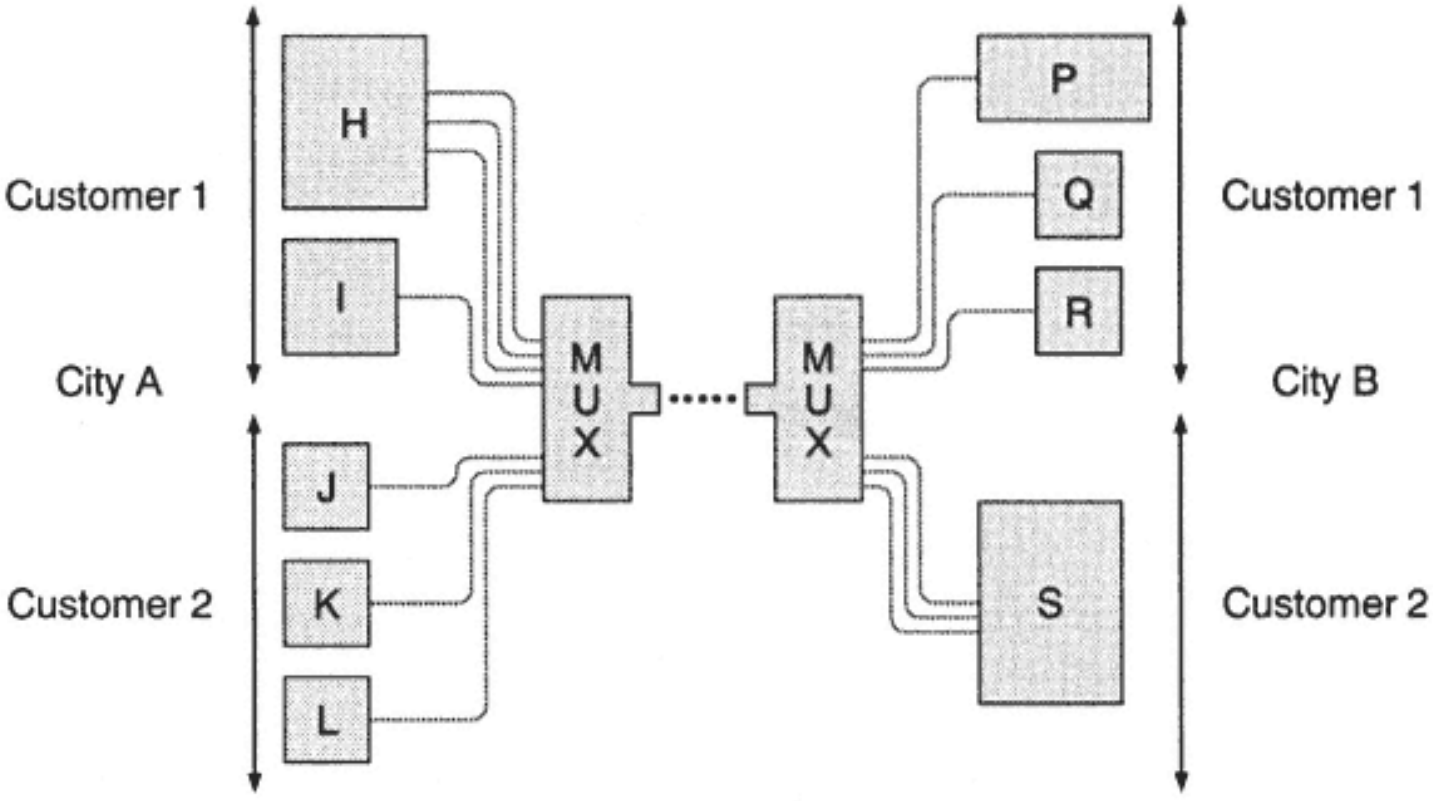

Figure 10.3
Sharing a MUX between different customers in different cities.

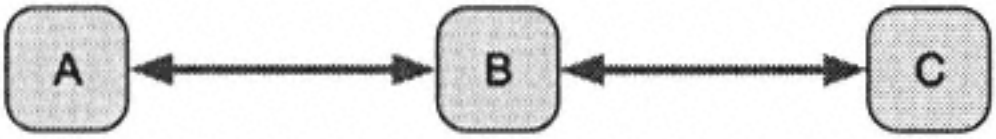

Figure 10.4
Two multiplexed channels connecting three cities.

tomer. However, the advantage of this approach is the flexibility it offers the user.

It may turn out that organizations in a third city (city C) also need to communicate with facilities in city B, resulting in another multiplexed channelthis time between cities B and C, as shown in Figure 10.4.

Now, what happens if a customer in city A wants to communicate with the group's office in city C? There may not be sufficient traffic to warrant another pair of multiplexers in cities A and C to connect the two. On the other hand, the costs of communicating between cities A and C directly may be so prohibitive as to force the abandonment of the planned communication.

Of course, there *is* a link between cities A and C (albeit an indirect one) via the multiplexers in city B. Communication between DTEs in cities A and C requires the interconnection of the two multiplexers in city B. Once that connection is made, communications between city A and city C can pass through city B, giving the impression that cities A and C are directly connected, as shown in Figure 10.5.

The interconnection between the two multiplexers in city B can be either

Direct. Each multiplexer treats the other one as one of its DTEs, as depicted in Figure 10.6. The functions performed by the multiplexers depend upon a number of issues. For example, if the connections are dedicated to specific DTEs in cities A and C, the multiplexers simply pass the data through the related connection. The connection between them is configured to give the illusion to MUX1 that MUX2 is the DTE in city C, and to MUX2 that MUX1 is the DTE in city A.

However, if the connections between the two multiplexers are accessible by any DTE, additional information must be passed between the DTEs prior to any information transfer. For example, if a DTE in city A is to send information to a DTE in city C, the following steps will probably be required:

1. Obtain a free channel connecting MUX1 to MUX2.

2. Inform MUX2 of the intended destination DTE in city C.

3. Allow the communication to proceed.

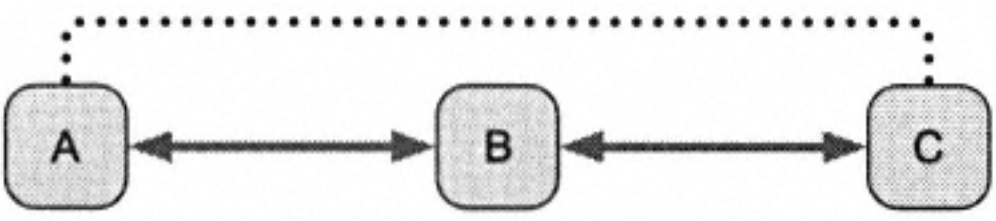

Figure 10.5
A logical connection between cities A and C.

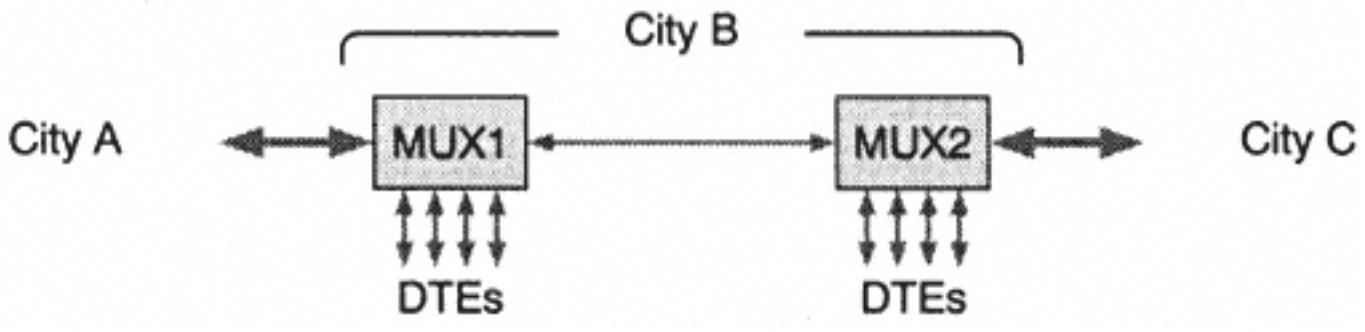

Figure 10.6
Connecting cities A and C by direct MUX interconnection.

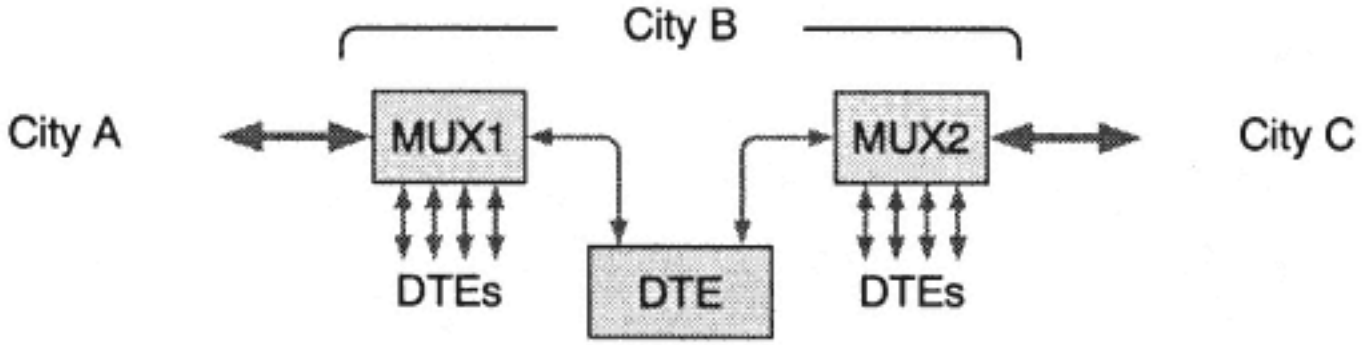

Figure 10.7
Indirect connection of MUXes by an intermediate DTE.

Indirect. Information is passed between a DTE connected to both MUXes, as shown in Figure 10.7. When data is sent to a DTE in a remote city (say, from city A to city C), it is first passed to a port on the DTE connected to both multiplexers. That DTE can then forward the data out of the port associated with the destination DTE in city C. The shared DTE can be made to support both communication methods described for directly connected multiplexers.

This approach offers considerably more flexibility than direct connection, since the tasks are *well defined*. That is, the multiplexers are responsible for handling multiplexed data and forwarding the data to a specific port, while the shared DTE is responsible for establishing the connections and taking the data from an incoming port and supplying it to another, outgoing, port.

In addition to the above, a subtle change has occurred in the way the communication takes place. Until now, all communications have been point-to-point with no intermediate devices other than those directly involved in the communication. Now, however, they can take place between a *number* of intermediaries.

10.3
Wide Area Networks

The previous section proposed the design of a system of interconnected multiplexers to support data traffic, rather than a combination of voice and data as in the existing telephone network. In addition, if statistical multiplexing is utilized, the bandwidth is employed only when necessary, potentially reducing the costs associated with using the telephone system.

In the mid-1960s, people working in the field of data communications recognized the potential of developing computer networks that would allow the transmission of data between geographically dispersed DTEs. Since these networks were envisaged to span whole continents (or even the entire world), they were called wide area networks, or WANs.

A wide area network is an interconnection of a number of multiplexing devices known as nodes. Each node in the network is directly connected to at least one other node through a communication channel known as a link. If a node is directly connected to *N* other nodes, it is associated with *N* links. Links are typically high-speed communication channels such as optical fiber, coaxial cable, microwave, or even satellite.

Additionally, each node supports a number of DTEs. A DTE can be directly connected to a node, or the node may support dial-in facilities through which remote users can access the wide area network. Once connected to the network, a DTE potentially has access to any other DTE on the network. A hypothetical wide area network is shown in Figure 10.8 (each square is a node, and the lines are links).

The internal structure of a node closely resembles the multiplexerDTEmultiplexer combination discussed in the previous section. A node performs a variety of functions, including

Allowing DTEs to send and receive information on the network.

Routing information to DTEs via other nodes. For example, in the wide area network in Figure 10.8, a DTE in Sydney could send information to a DTE in Paris via the Tokyo, New York, London (and possibly Rome) nodes.

Multiplexing information on the links, thereby sharing the links among a number of DTEs. For example, all the DTEs attached to the Sydney node would share the link to Tokyo in order to reach DTEs on nodes other than Sydney.

How these functions are realized in a wide area network depends upon a number of factors, some of which are now considered.

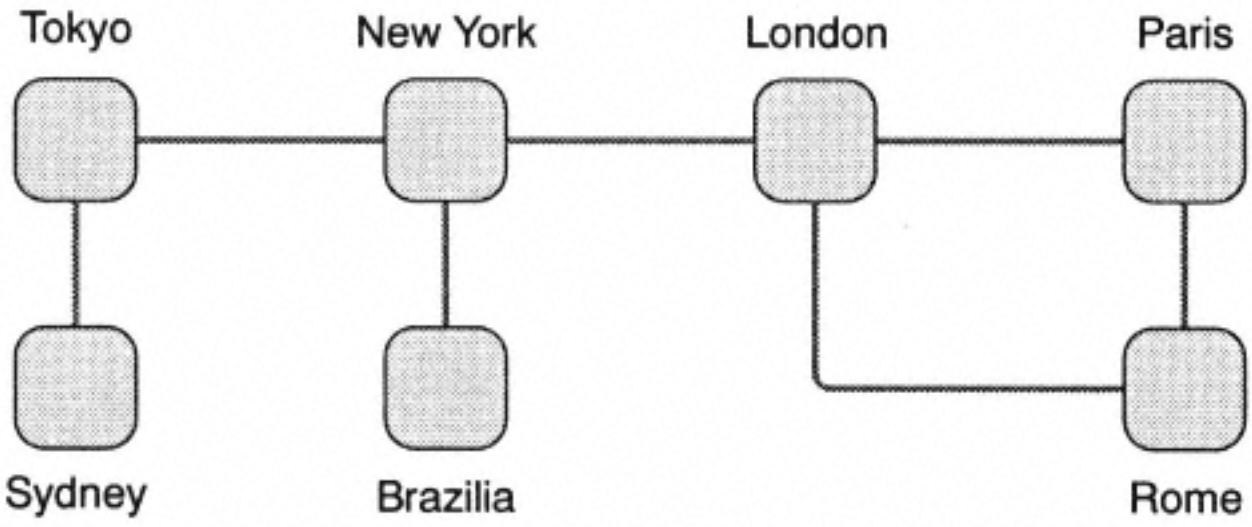

Figure 10.8
A hypothetical wide area network.

10.3.1
Topologies

There is no prescribed "shape" to a wide area network, since the geographic area covered may vary from country to country or continent to continent. However, if the geographic factors are ignored, there are a limited number of topologies:

Fully connected. Every node on the network connects to every other node. If the network consists of *N* nodes, each node will require *N* 1 links. For example, the network in Figure 10.9 has six nodes, and each has five links.

Fully connected networks are usually very costly because of the number of links required. However, they are usually very fast (since there are no intermediate nodes to handle a message), and they offer a high degree of reliability (if a link fails, there should be an alternate path to the intended destination).

Star. All nodes connect to a single, central node (often referred to as the hub); all communications between nodes pass through the hub node. Figure 10.10 is an example of a star network.

Although the star reduces the number of links to a minimum (*N* nodes require *N* 1 total links), the failure of the hub node will result in the complete cessation of all communications.

Mesh. Probably the most common of all wide area networks. Usually distance or known traffic volumes determine which of the nodes are connected; a possible mesh network is shown in Figure 10.11.

Mesh networks can also evolve as communication patterns change over time. For example, a city might not initially have a node, but if it subsequently gains political or economic stature, one might be added. Similarly,

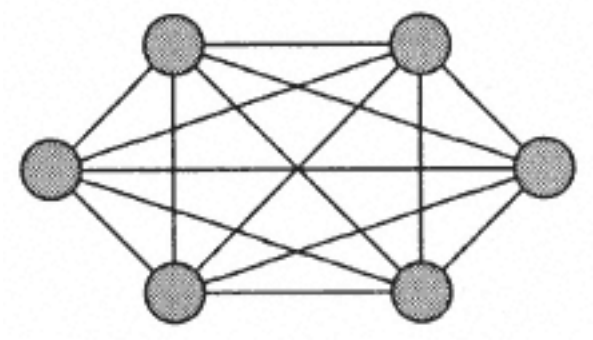

Figure 10.9
A fully connected network.

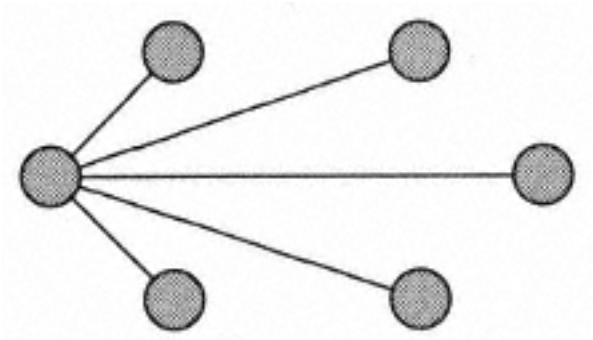

Figure 10.10
A star network.

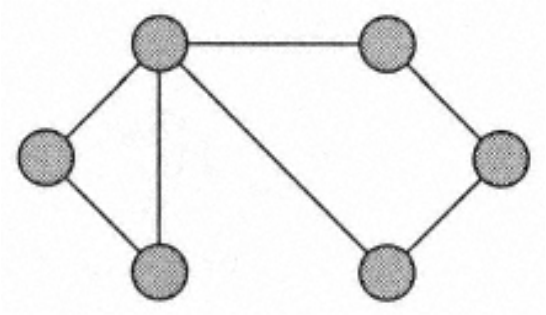

Figure 10.11
A mesh network.

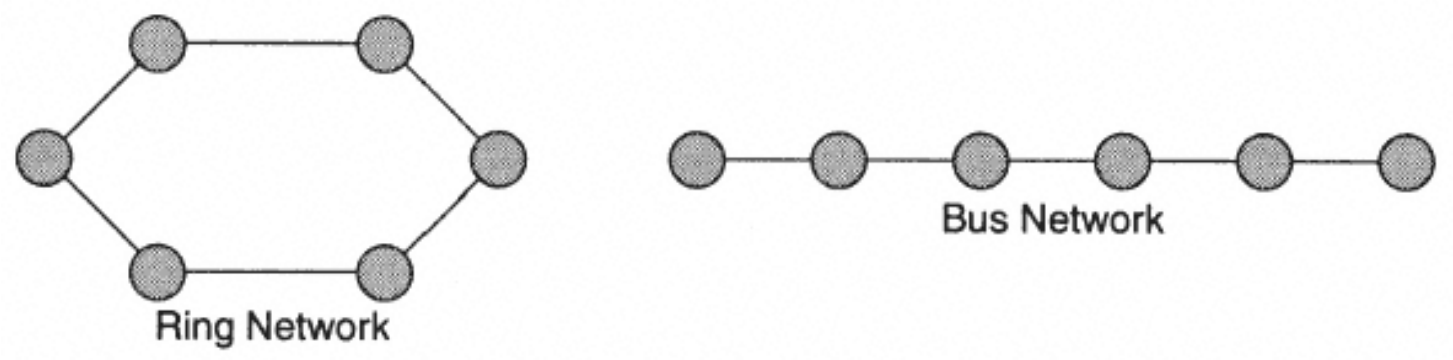

Figure 10.12
The ring and bus network topologies.

the traffic volumes between two cities might be low, justifying the use of an intermediate node; however, over time the volumes might grow to the point where a separate link directly connecting the two cities may be needed.

Other topologies. A number of other network topologies are possible. Two that will be discussed later in the book are the ring and the bus, shown in Figure 10.12.

10.3.2
Addressing DTEs

Sending information on a wide area network and sending information in a point-to-point communication are not the same. In a point-to-point communication there is only one possible destination (the DTE at the other end of the channel), whereas in a wide area network there can be any number of possible destinations. Therefore, whenever a DTE supplies information to a node for transmission, the node must know, at a minimum, the intended destination.

The intended destination is simply another DTE attached to a node somewhere on the network. How much information does the node require in order to get the information to the destination? Consider how a letter addressed to someone in another country is processed by the post office. The one thing the post office doesn't have is a list of all the people in the world, but it does have a list of countries and, possibly, regions within countries. Thus, the letter is first sent to the country in question, then to a region within the country, then to the city, and so on, down to the recipient's house. In other words, a hierarchy is used to route the letter.

The same approach can be applied to sending information in a wide area network: the information is not sent directly to a DTE or a port on a DTE. Instead, it is sent to the node to which the destination DTE is attached. The node can then supply the information to the DTE or perhaps a port on the DTE (assuming that the DTE supports several connections to the node).

All networks require some form of address to ensure that the information arrives at the intended destination. Broadly speaking, there are two types of address:

Hierarchical. This is a single address consisting of a series of fields, each of which is used to identify part of the destinationfor example, the node, the DTE, and the port on the DTE.

Global. This is a single address that uniquely identifies a single entity on the network (typically a DTE attached to a network node).

All DTEs attached to the network are associated with an address, meaning that potentially any pair of DTEs can communicate, assuming they agree to the same protocol.

In addition to the destination address, most, if not all, wide area networks require the address of the transmitting DTE (the source address) to be sent. The source address can be used for a number of different purposes: the most obvious is to allow the receiving entity a means whereby a reply or an acknowledgment can be returned.

10.3.3
Information Representation

To be effective and to attract as large a market as possible, a wide area network often supports a variety of protocols, thereby allowing different DTEs access to it. For example, the network should permit asynchronous DTEs to co-exist with synchronous DTEs. In addition, it should be transparent: DTEs accessing the network should not be required to change their protocol. Thus, if a telephone circuit between a pair of asynchronous DTEs is replaced by a wide area network, the asynchronous communications should continue as before.

These two seemingly contradictory objectives (co-existence of protocols and transparency) require the wide area network to support a universal data structure that is internal to it and recognized by the nodes, regardless of the DTE information carried. Commonly known as a packet, the structure should be able to carry the information specific to the protocol as well as control information specific to the wide area network. Typical control information can include the source and destination addresses as well as some form of error checking. A generalized packet structure is shown in Figure 10.13.

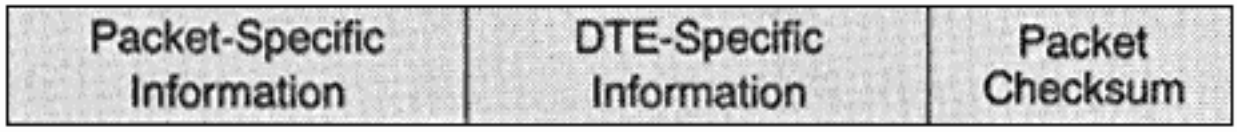

Figure 10.13
A generalized packet structure.

10.3.4
Packet Routing

The various nodes are responsible for the routing of packets through the network, from the source node to the destination node and eventually the destination DTE. There are a number of similarities between a circuit-switched telephone call and packet routing; accordingly, wide area networks are often labeled packet switching networks, or PSNs. (The terms Packet Switching Exchange (PSE) and Packet Switch Node (PSN) are synonymous with *node*.)

The algorithms chosen to support packet routing in a network depend upon factors such as

The cost of sending a packet across a specific link. Some links may be more costly than others: terrestrial links are typically cheaper than satellite links.

The length of time it will take the packet to reach the intended destination. Time-critical information may be required to move through the network with a higher priority than other information.

The volume of traffic on the network. As the number of packets on the network increases, congestion can occur, leading to delays and other problems (see below).

The order in which the packets will arrive at the intended destination. If multiple paths exist in the network and packets are permitted to take any path, then those taking a longer path may arrive at the intended destination later than other packets, even though the late packets were sent earlier.

How lost or duplicate packets are handled. In any communication, information can be lost through some form of error. In a wide area network, mechanisms may be required to ensure that information always reaches the intended destination. Similarly, if multiple paths exist and a packet somehow is sent on two or more links, mechanisms may also be required to recognize the duplication.

When a packet is transmitted through a wide area network, it will be handled in one of three ways:

1. The node attached to the source DTE will accept the DTE's information, store it in a packet, and transmit the packet onto the network.

2. The node attached to the destination DTE will remove the packet from the network. The information in the packet will be transmitted to the destination DTE using the DTE's protocol.

3. Any intermediate nodes will route the packet to the final, destination node using some type of routing strategy or routing algorithm.

Routing Algorithms

Many different algorithms can support packet routing through a network. At one extreme, the node can simply forward the packet through all of its links in the hope that a copy of the packet will eventually reach the intended destination node. This is known as flooding, and although there is a very good chance that one copy of the packet will reach the destination, there is also a very good chance that many other copies will as well. In addition, copies of the packet could be traveling through the network for indefinite periods of time before reaching the destination (this is especially true if there are loops or cycles in the network). Two ways of improving this algorithm are

Supplying each packet with a hop countor counter, that is given an initial value by the source node and is decremented by each node that receives the packet. When the hop count reaches 0, the receiving node discards the packet. The hop count must be large enough to ensure that nodes at opposite ends of the network can communicate.

Transmitting the packet on all links *except* the one on which it was received. This helps to reduce the number of unnecessary transmissions between pairs of nodes. The assumption behind this algorithm is that if a node forwards a packet, it wasn't the intended destination; therefore, there is no need to retransmit the packet back to the original node.

There are a number of limitations to the above approaches, chief among them the unnecessary routing of packets to nodes that aren't on the path to the intended destination.

An alternative to flooding and its variations is to supply each node with a data structure known as a routing table. In its simplest form, the routing table contains a list of all possible destination nodes and the link to the next node on the path to the specified destination. For example, using the network illustrated in Figure 10.14, station A's routing table is as depicted in Table 10.1.

In Table 10.1, the destination is the intended destination of the packet supplied to the node and the next node indicates the link to be taken to reach the destination (a dash signals that there is no next node). For example, to reach nodes B, D, E, and F from node A, all packets would take the link to

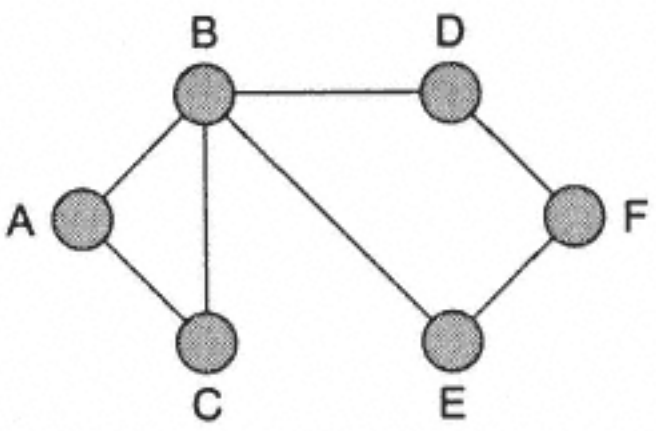

Figure 10.14
A wide area network of six stations.

Table 10.1 Routing table for station A.

Destination	A	B	C	D	E	F
Next Node		B	C	B	B	B

node B; however, to reach C, they would be sent directly on the link to node C.

The choice of routes becomes somewhat more complex when considering node F. The routes from F to nodes D and E are readily apparent, but there are two choices from F to B (via either node D or node E).

There are many ways in which the routing table can be constructed, organized, and used; some of which include the following:

If the traffic patterns on the network are well known and the topology is stable (i.e., the addition or removal of a node or link is a rare event), it may be possible to decide upon the routes when the network is initially implemented. These routes can be placed in the routing tables of the different nodes and the network left to function. This is often referred to as static allocation.

For example, it may be decided that the route from node F to node B is via node E, producing the routing table for node F shown in Table 10.2. The choice of routes can be based upon any number of factors, typically including the cost of transmitting packets across a link, traffic patterns (some nodes may be subject to delays), and distance.

If there are multiple paths between nodes, the routing table can contain information about them; this is known as multipath routing. For example, node F's routing table could be expanded to show that a route to node B also exists via node D (in addition to the route through node E). Multipath routing has a number of uses including: traffic can be shared between the different routes, and should a route become inactive because of a link or node failure, an alternate route may be available.

In many situations, the topology and traffic patterns of the network can changefor example, because of node or link failures and variable traffic patterns during certain times of day. If node E fails, traffic from node F may have to be rerouted through D to reach B. Algorithms that allow the modification of routing tables are known as dynamic or adaptive routing strategies.

A variation on the use of static allocation is to have a network control center that periodically sends new routing tables to the nodes. These ta-

Table 10.2 Routing table for station F.

Destination	A	B	C	D	E	F
Next Node	E	E	E	B	B	

Table 10.3 Routing table for node D. (Link BD not functioning.)

Destination	A	B	C	D	E	F
Next Node	F	F	F		F	F
Count	4	3	4		2	1

bles are based upon information that the network control center receives from all of the nodes; each node sends packets containing network status information. For example, the network control center could be informed of node E's failure and could then produce new routing tables that would have node F's traffic flow via node D.

The use of a centralized network control center to perform routing table updates can be costly to the network in that a percentage of the network bandwidth is lost to the exchange of status information. Ideally, the amount of bandwidth lost will be minimal. One way to minimize bandwidth loss is known as backward learning, in which every packet contains a count that is increased as the packet is transmitted by each node (the further away a packet is from its source, the greater the value of the count). Each routing table entry has a distance count associated with every destination in the table (the count indicates the number of nodes the destination is away from this node). Initially, these distance counts are set to infinity.

The backward learning algorithm works as follows. Whenever a packet is received by a node, the destination entry in the routing table that corresponds with the *source* address in the packet is found. The routing table count is compared with the packet's count; if the packet count is less, the routing table entry is changed to the packet count; otherwise, nothing is done.

For example, if the link between B and D is not functioning, the routing table for node D might look like that shown in Table 10.3. Should the BD link become active, a packet from B might arrive across the link at D with a count of 1. Since this count is lower than the existing routing table count of 3 for destination B, the next node for destination B changes to B and the count changes to 1. This is shown in Table 10.4. As soon as a packet from D arrives at B, B's routing table count for D changes, meaning that messages from A and C can take the shorter route.

Table 10.4 Routing table for node D. (Link BD not functioning.)

Destination	A	B	C	D	E	F
Next Node	F	B	F		F	F
Count	4	1	4		2	1

10.3.5
Methods of Packet Transmission

Packets are transmitted across a network one node at a time. Ideally, the packet will reach the intended destination with the minimum of difficulty, but over a period of time, nodes may fail (i.e., crash because of a software or hardware fault) and links may fail (e.g., being cut by someone digging a trench). Should either of these types of fault occur, the packet transmitted might not reach its destination. The remedial action depends upon the type of service being offered by the network. Three common methods of service are now considered.

Datagram

A datagram service is a *best-effort* communication service that does not guarantee anything about the transmission of the packet. For example, a packet may be sent and then lost because of a network fault, or the packet might simply be discarded as a result of heavy network traffic flow (i.e., network congestion). On the other hand, duplicate copies of the packet might arrive at the destination because of routing problems in the network or simply the routing algorithm itself (consider the possible effects of flooding).

In a datagram service, it is assumed that the DTEs are responsible for handling any errors associated with datagram transmission, perhaps by constructing their own go-back-N protocol atop the datagram service.

Messaging

A messaging system is one in which a message is sent as a series of packets that are reassembled into the complete message at each node before being forwarded to the next link.

In a typical messaging application, a DTE divides a large set of data, the message (such as a file), into a series of packets, which are transmitted between nodes. The receiving node reassembles the message from the incoming packets, and then performs its routing algorithm and determines the next link. The message might not be transmitted immediately; the node may perform some type of scheduling based on criteria such as the amount of buffer space available on the next node and the amount of traffic on a particular link. Networks that support messaging are also referred to as store-and-forward networks.

Messaging systems normally offer "guaranteed" service, in which a message sent from one DTE will eventually arrive at the intended destination DTE. To achieve this level of service, the communications *between* each node must ensure that each packet has been successfully transmitted. Protocols similar to Bisync or SDLC are often used by the individual nodes when communicating across the link.

Messaging is often used in situations where interactive computing is not required and the arrival of the data at the destination is not time critical.

Virtual Circuits

Some applications cannot function using datagrams or messaging. Consider, for example, an interactive terminal connected via a wide area network to a remote host. The user of the terminal wants to make sure that the data entered arrives at the remote host when and in the order it is entered. Imagine the fun of trying to edit a file in which each line is sent as a datagram: some lines might never arrive, and others could arrive several times.

Network designers have recognized these limitations and have set about developing a system analogous to the public telephone network's "switched circuit" known as a virtual circuit. This is a path through a wide area network that is maintained for the duration of the communication (as is a switched circuit in the telephone network). The data entered by the user follows the same path through the circuit until the communication is ended. If another communication to the same destination is started at a later time, its path can differ from the previous one because of traffic conditions and changing topology. Virtual circuits are reliable: the data supplied by a transmitting DTE arrives in the order sent, without loss, duplication, or error, at the receiving DTE. Over a period of time, there can be hundreds of simultaneous virtual circuits established across the network.

A virtual circuit consists of three distinct phases:

Call establishment. The calling party (a DTE) must indicate to its node that a virtual circuit is requested; a request includes the address of the called party (normally a DTE attached to another node). This information is put into a call request packet that is routed through the network by the various nodes on the path between the two DTEs. The network's routing algorithm dictates the path taken by the call request packet.

As the call request packet progresses through the network, a logical path is established between the nodes on the network for this specific call. Each node is responsible for maintaining its part of the virtual circuit for the duration of the call in its virtual circuit table. When the call request packet finally arrives at the destination node, the destination DTE is "called," and if the call can be accepted, a call connected packet is returned to the calling DTE. However, if the call cannot be accepted (for example, if the DTE is inactive), it will be cleared by the called node returning a call clearing packet to the calling DTE.

Although a virtual circuit itself is full-duplex, the path created by the call request packet need not be. For example, the call request packet could set up a simplex channel, in which case the call connected packet establishes a return simplex channel. It is possible for the two halves of the virtual circuit to take completely different paths through the network.

Data transfer. Once the call has been established, data transfer can take place. Data sent by either DTE passes through the network in a data packet using the virtual circuit created in the call establishment phase.

Since there can be any number of virtual circuits in use, each packet is sent with an indication of its virtual circuit, the virtual circuit number. As the data packet arrives at each node, the virtual circuit number, in conjunction with the virtual circuit table, indicates the next leg of its journey.

To ensure the reliability of the data transfer, data packets are acknowledged, typically as they move between nodes. In some cases, there can be end-to-end acknowledgments, confirming the arrival of the packet at the destination node.

Call clearing. As in any communication, there comes a time for the connection between the two DTEs to be terminated. Since there are a finite number of possible virtual circuits in the network (there is a limit to the size of the virtual circuit tables), both the DTEs *and* the nodes involved in a virtual circuit must be informed of the call's termination.

A call is terminated by either DTE sending a call clearing packet. The packet passes through the original path created during the call establishment phase, and each node is responsible for freeing its part of the virtual circuit. The normal response to a call clearing packet is a confirmation packet.

Although it may appear that a single virtual circuit number is used to identify a virtual circuit through the network, this is seldom the case since the overheads associated with maintaining globally unique virtual circuit numbers outweigh their benefits. Instead, what appears to be a single virtual circuit is usually made up of a number of separate virtual circuits, allocated on a node-by-node basis. Consider the virtual circuits shown in Figure 10.15 that utilize the links connecting nodes X, Y, and Z, in which virtual circuits exist for DTEs A and D, as well as DTEs B and C.

In the above example, process B communicates with process C through node X's virtual circuit number 2, whereas C uses virtual circuit number 1 on node Z. Since the virtual circuit numbers change from link to link, the node is responsible for changing the packet's virtual circuit number prior to being transmitted out the next link. Consider a packet sent from B to C: the initial

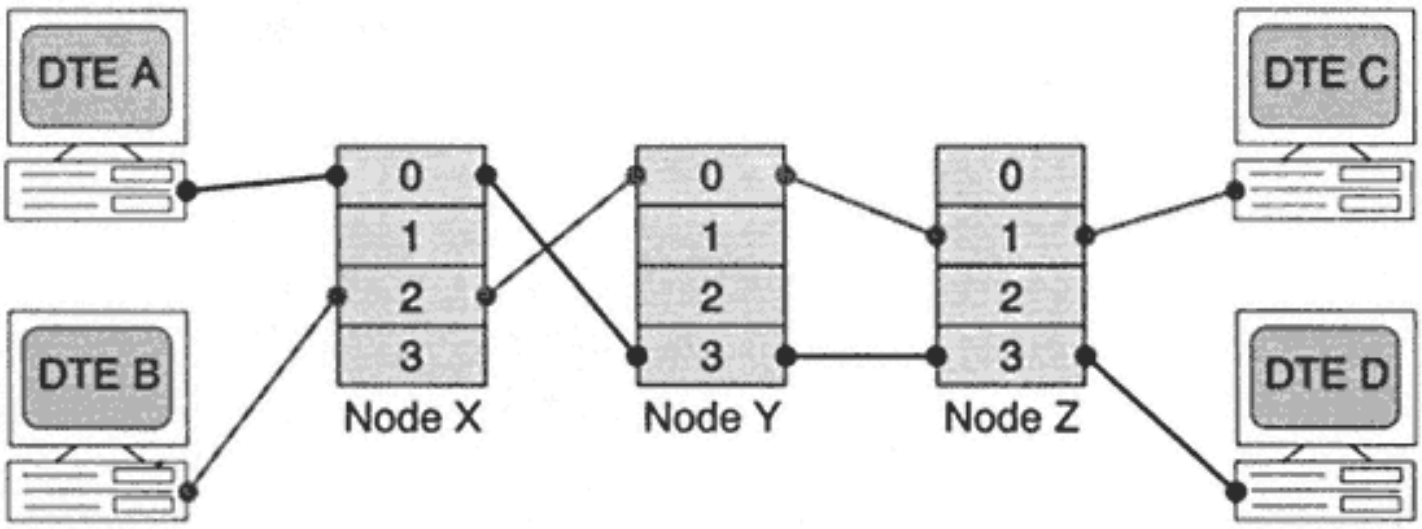

Figure 10.15
Virtual circuits connecting DTE A
to DTE D and DTE B to DTE C.

virtual circuit number is 2, which is changed to 0 when the packet is sent to node Y; at node Y the number is changed again, to 1, and then sent to node Z, which forwards the packet to process C.

To distinguish between datagrams and virtual circuits, datagram services are often called connectionless and virtual circuits are called connection-oriented.

Many wide area networks offer combinations of the above services. For example, virtual circuits for interactive users may be needed during working hours, but during the quiescent periods at night, large-volume transfers can be sent using a messaging service.

10.3.6
Node Design

The exact functions of a node depend upon a number of factors, such as the protocol used by the wide area network and the protocols supported by the DTEs. However, there are a number of generic features that a node should support:

The transmission and reception of packets from other nodes.

The routing of packets from one node to the others.

The control of the DTEs' network access.

Although it is possible to design a node so that all of these functions are performed in one large maze of software, network designers have, over the years, come to recognize that dividing node software into layers or distinct parts offers a number of benefits, including

Faults are easier to isolate and fix.

Software testing is simplified.

One layer can be modified while isolated from the other layers.

The most natural (and perhaps obvious) division of the node is between the part that controls the network and the part that allows the DTEs network access. The network part is responsible for forwarding packets and maintaining the virtual circuits; the DTE part allows DTEs to send and receive information across the wide area network. The entire structure is depicted in Figure 10.16.

Within each of these two broad divisions, further layering is possible.

Network Software Layering

Consider the functions associated with the transmission and reception of packets in a network supporting virtual circuits or messaging, described below.

Nodes are connected to other nodes by links. A single node may be associated with a number of links, not all of which will support the same

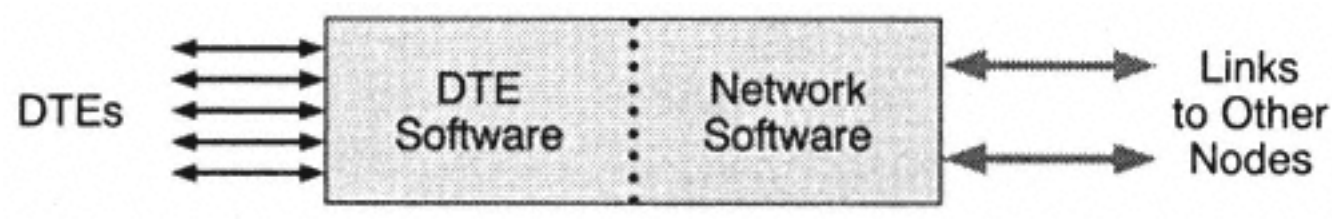

Figure 10.16
Functional divisions of a node.

medium (for example, some may be high-speed synchronous lines, others may be asynchronous, and still others may be satellite). In many wide area networks, the lowest level is known as the Physical Layer and defines the electrical and mechanical properties of the link. The nodes at the end of each link must agree to the same Physical Layer protocol.

Both virtual circuits and messaging require that packets sent across a link arrive undamaged and in the correct sequence at the receiving node. Since the underlying Physical Layer offers only a transmission medium, another level, the Data Link Layer, is required, supporting the following features:

A device driver, specific to the underlying medium, that is responsible for the transmission and reception of packets.

A reliable, node-to-node protocol, such as stop-and-wait or go-back-N. This feature ensures that any packet sent across the link either will arrive correctly or will be received in error and a retransmission requested. (In a datagram network, the node-to-node protocol need not be reliable.)

The Data Link Layer is intended to ensure that packets sent across the underlying link arrive at the next node in the correct order and error-free. Although the device driver is unique to each type of link, the reliable protocol software need not be. That is, packets are formatted with the correct sequence number and control fields by the protocol software, while the underlying device driver and the Physical Layer simply transmit the information supplied. Once again, the argument for using a single set of software deals primarily with maintenance: with one set, "fixes" are universal.

Finally, calls must be established and packets routed through the network. Since these functions cannot be justifiably performed by either the Physical Layer or the Data Link Layer, a third level, the Network Layer, is introduced. This layer is responsible for a number of functions, including

Inspecting packets received by a Data Link Layer to determine the destination of the packet. The packet may be intended for a DTE attached to the local node, in which case the Network Layer supplies the information in the packet to the DTE. Or it may be destined for a remote DTE, requiring the Network Layer to forward it to the Data Link Layer that will eventually lead to the destination DTE.

Taking information from the DTE (typically data and the address of the intended destination) and passing it in a packet to the correct Data Link Layer for eventual transmission. Call request and call clearing

packets are also handled by the Network Layer and forwarded to a Data Link Layer for transmission.

Should the destination DTE be attached to the same node as the transmitting DTE, the Network Layer need not attempt to transmit the information; instead, the information can be supplied directly to the specified DTE. The Network Layer acts as a multiplexer both to the DTEs that are attached to the node and to the links.

The relationship between the various layers, the DTEs, and the connections to other nodes is shown in Figure 10.17.

Layering involves more than simply dividing tasks into logical functions: it encompasses data structures and, in some respects, programming style. Many layered communication systems define a single layer structure and apply this structure to all layers in the system:

Each layer has its own protocol and protocol conventions. Communications occur between peer entities *within* a layer. For example, a wide area network's Network Layer has a set of protocols specific to it. These protocols deal with the creation, maintenance, and termination of virtual circuits (or perhaps simply the routing of datagrams). In the Data Link Layer, protocols might exist for the reliable transfer of information between nodes.

Communications within a layer occur between peer entities, however, these entities use the communication services offered by the layer below. At a minimum, the services associated with a layer consist of the transmission and reception of information. The communication system may offer features that permit a lower layer to return error status information to the higher layer.

Since each layer has its own protocol and uses the services of the underlying layer, it is necessary to encapsulate one layer's packet structure in the packet structure of the underlying layer. Encapsulation involves adding

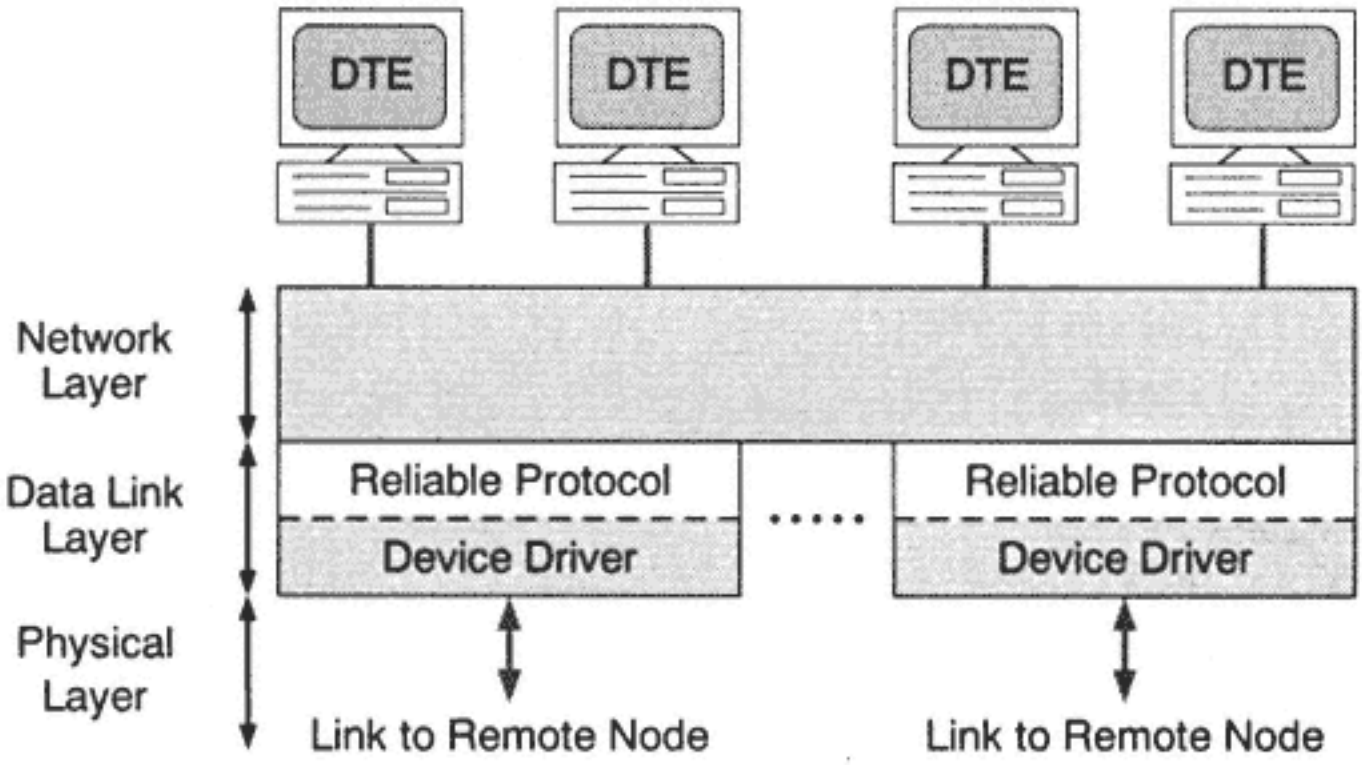

Figure 10.17
A node's layers and their relationships.

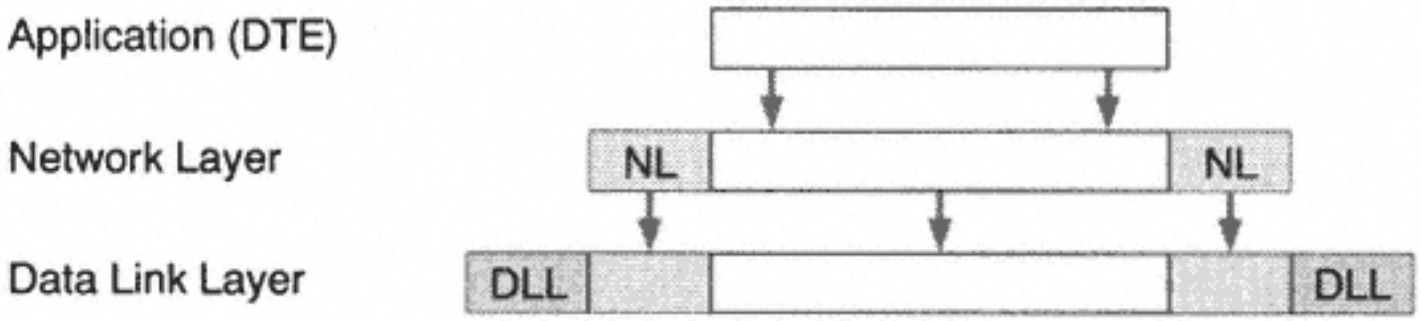

Figure 10.18
An example of encapsulation.

control information specific to a layer (such as sequence numbering and error checking). Thus, a packet sent by a DTE in the layered system depicted in Figure 10.17 could proceed through the sequence of encapsulation steps shown in Figure 10.18. Finally, the data within the packet is subject to the rules associated with transmitting information in the Physical Layer. For example, bit stuffing or *DLE* insertion could occur. When the packet is received, the control information associated with a specific layer is removed before the packet is passed to the layer above.

The path of a DTE's packet through a wide area network may take it through many intermediate nodes. Figure 10.19 shows how an application on a DTE attached to node A communicates with an application attached to node C through an intermediate node, B. A packet sent by the application attached to node A is encapsulated by each layer on node A and eventually sent to node B. The packet arrives at node B's network layer, which is responsible for routing it by consulting either a routing table or a virtual circuit table. In either case, the packet is re-encapsulated and sent out the link attached to node C. At node C, the network layer supplies the packet to the destination application.

The Node-DTE Interface

The node and its DTEs must interact when communications across the wide area network are to take place. Layering still applies between each DTE and the node. In this case, the DTE must inform the Network Layer of the action to be taken with the information suppliedfor example:

DTE to node. The DTE can request a call to be made to another DTE on the network; data can be sent and received from a remote DTE; and the DTE can request that the call be cleared.

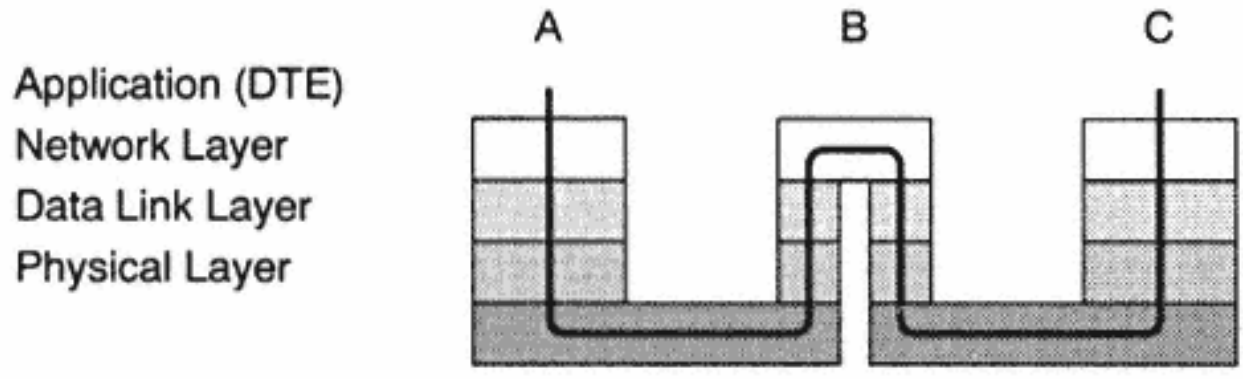

Figure 10.19
The path taken by a packet from node A to node C via node B.

Node to DTE. The node can supply the DTE with requests for setting up a call, data from a remote node can be supplied to the DTE, and the node can indicate to the DTE that the call is being cleared.

The separation of the DTE from the node is logical but somewhat arbitrary, since there are many situations where a DTE cannot communicate directly with a node and other situations in which the node and the DTE are indistinguishable. Consider the following:

DTE and node combined. In an attempt to reduce costs, many networks have node software that is intended to function on interconnected DTEs (i.e., the DTEs themselves are the nodes). The DTEs communicate via leased lines or, in some cases, dial-out facilities. Packets are formatted by software (typically within the operating system) and forwarded to the Network Layer. In this design, not only are costs reduced, but the features offered by the operating system (for example, file storage for messaging) can simplify the design of node software.

DTE and node separate. In most commercial wide area networks, the node and the DTE are separate and distinct. If the DTE is programmable, packets can be formatted by the DTE according to the Network Layer's requirements and sent to the node for transmission. However, if the DTE is simply an interactive terminal (or a computer emulating a terminal), a number of problems arise; for example, where is the data to be formatted into a packet, and how does the node determine the intended destination of a packet?

The solution adopted by many wide area networks is to develop a software package known as a Packet Assembler/Disassembler (PAD) that on the DTE side emulates a host to which the DTE communicates, and on the node side generates packets based upon the Network Layer's requirements (the node can be connected to other DTEs and PADs). As shown in Figure 10.20, an asynchronous terminal could communicate with a PAD, which encodes and decodes Network Layer packets.

An alternative to having a separate PAD is shown in Figure 10.21, in which the PAD is part of the node. The number of DTEs that could connect to an individual PAD would be implementation dependent.

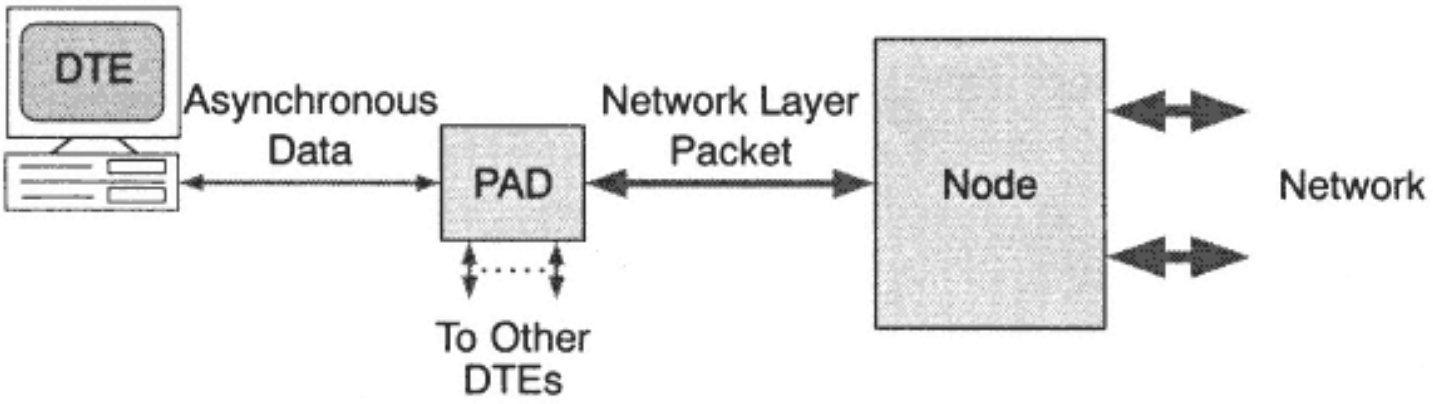

Figure 10.20
Connecting to a network via a PAD.

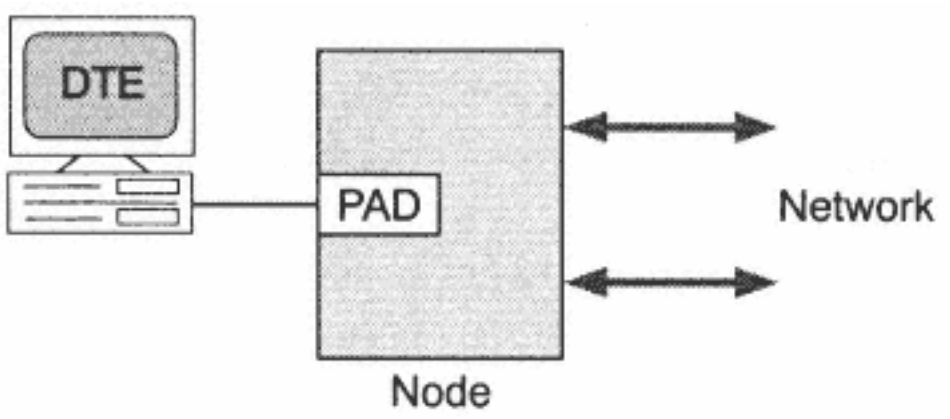

Figure 10.21
A PAD can be part of a node.

PADs are not restricted to asynchronous DTE communications; they are available also to support Bisync and other such protocols.

10.4
Wide Area Network Protocols

From the discussion in the previous section, it should be apparent that there is no single wide area network protocol: each layer has its own set of protocols and standards. In some cases, there can be numerous choices of protocol within a single layer.

10.4.1
Physical Layer Standards

The Physical Layer connects pairs of nodes through a communication channel. Physical Layer standards define the mechanical, electrical, functional, and procedural rules associated with the communication channel. Examples of Physical Layer standards include RS-232-C, RS-449 (RS-423-A and RS422-A), X.21, and X.21*bis*.

10.4.2
Data Link Layer Protocols

The Data Link Layer is responsible for formatting information for transmission across the underlying Physical Layer. Typically, this layer offers a reliable communication service to the higher layers by maintaining error detection and flow control. A network supporting datagrams only does not have the same need for reliability.

There are a number of different Data Link Layer protocols. Some networks use Bisync or SDLC, neither of which is considered a standard. Internationally adopted standards, all based upon SDLC (see Chapter 6), include

ADCCP. The Advanced Data Communication Control Protocol, developed by the American National Standards Institute (ANSI) as ANSI standard X.366.

HDLC. The High-Level Data Link Control protocol, a point-to-point synchronous protocol intended for Data Link Layer peer communications, developed by the International Organization for Standardization (ISO). HDLC is defined in ISO documents DIS 3309 and DIS 4335.

LAP-B. The Link Access Procedure-Balanced protocol is defined by ITU-T as part of the X.25 protocol (see below). LAP-B is a subset of HDLC.

10.4.3
Network Layer Protocols

Probably the best known Network Layer protocol is X.25, an ITU-T standard developed to handle virtual circuits between DTEs and a wide area network. In X.25, the node is referred to as the DSE (Data Switching Exchange), and the connection between a DTE and a DSE is the DCE (data circuit *terminating* equipment). The protocols used within the network are not defined by X.25.

X.25 is actually an ITU-T *recommendation* as opposed to a specific protocol, consisting of a detailed description of the Physical, Data Link, and Network Layers:

Level 1. The physical, electrical, functional, and procedural characteristics needed to establish, maintain, and disconnect the physical link between the DTE and the DCE. The DTE/DCE interface characteristics are defined in Recommendations X.21 and X.21*bis*.

Level 2. The link access procedure for data interchange across the link between the DTE and the DCE. Level 2 access is defined by the LAP-B protocol.

Level 3. The packet format and control procedures for the exchange of packets containing control information and user data between the DTE and the DCE. Level 3 is also referred to as the Packet Level DTE/DCE Interface and describes X.25 virtual circuits. Many public data networks have adopted the X.25 level 3 protocol for network access.

The structure of X.25 is shown in Figure 10.22.

Level 3 of X.25 defines two types of virtual circuit: virtual call (also known as switched virtual circuit, or SVC) and permanent virtual circuit, or PVC. Switched virtual circuits were described in previous sections. Permanent virtual circuits are those in which the virtual circuit is always available for

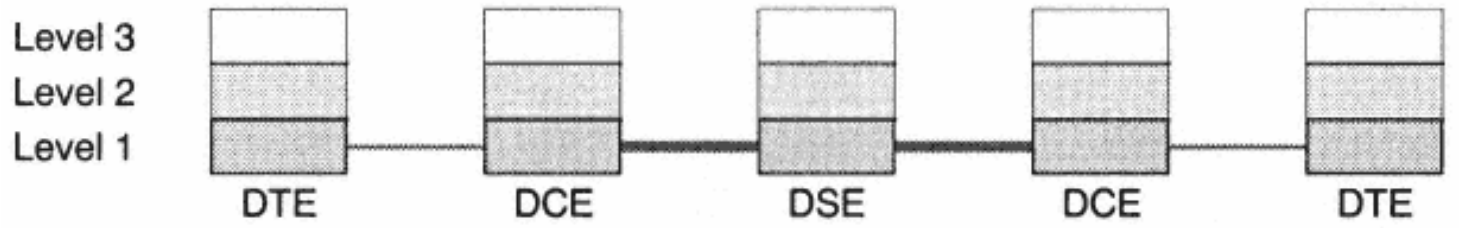

Figure 10.22
The structure of X.25.

data transfer. There is no call establishment or call clearing phase associated with a permanent virtual circuit.

The three phases associated with an X.25 virtual circuit are described in the following paragraphs.

In the call establishment phase, a call can be established by any DTE that has an available virtual circuit number. The call must be sent in a Call Request packet, the format of which is shown in Figure 10.23 (note that ITUT labels its bits from left to right, with bit 7 being the low-order bit; an 8-bit quantity is referred to as an octet).

The various fields in the Call Request packet are defined as follows:

MOD. Informs the called DTE of which data packet sequence numbering scheme the calling DTE is using: either 8 or 128 (see "Data transfer," below). If the two bits are 01, modulus-8 is used; however, if the value of MOD is 10, modulus-128 is used. The two other possible bit patterns are not defined. All packets are sent with the MOD identifier.

Group/channel. A 12-bit virtual circuit number made from a 4-bit group number and an 8-bit channel number. All packets are sent with the 12-bit virtual circuit number.

Calling-len and called-len. The calling DTE's address length (calling-len) and called DTE's address length (called-len). Addresses can be up to 14 digits long.

Calling and called address. The source and destination addresses. The network address format is defined by ITU-T X.121: three digits for the country, one for the network within the country, and ten to identify the DTE. Each digit is four bits long (i.e., a binary-coded decimal digit). Within a network, the organization of the ten digits is defined by the network administration. If the final octet contains a single digit (i.e., 4 bits), the remaining four bits are 0-filled.

Fac-Len. The length, in octets, of the facilities field.

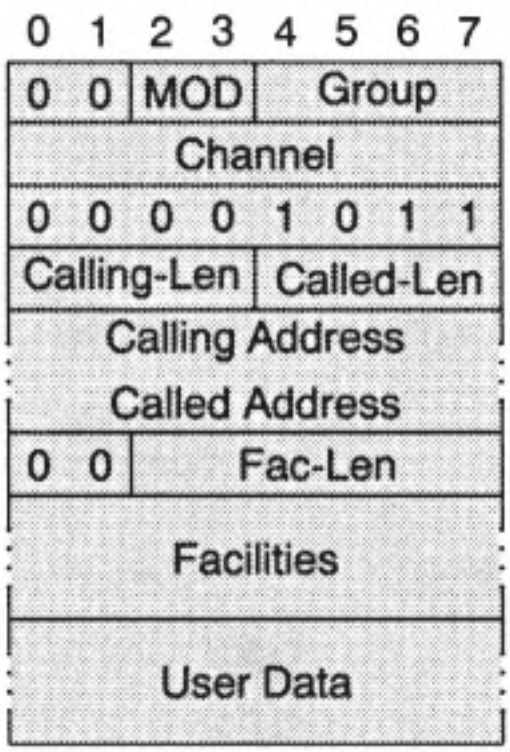

Figure 10.23
The X.25 Call Request packet.

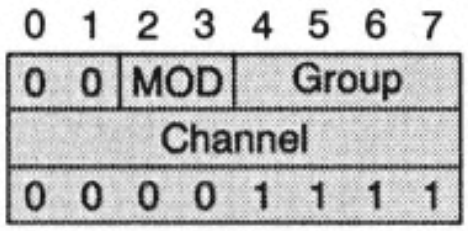

Figure 10.24
The X.25 Call Connected/Call Accepted packet structure.

Facilities. Allows the calling DTE to either indicate or request one or more facilities. The exact facilities supported can vary from network to network, but some of the more common ones are reverse charging (the calling DTE wants the called DTE to accept the network charges for this call), use of a nonstandard window or message size, and whether the calling DTE is a member of a specific Closed-User Group, or CUG (this is a facility that restricts communications to a specific group of DTEs). Facilities are sent as sequences of octet pairs: the first octet, the facility code, indicates the facility requested, and the second octet, the facility parameter, is the value associated with the requested code. For example, to request reverse charging, the facility code is 00000001, while the parameter is either 00000000 (no reverse charging requested) or 00000001 (reverse charging requested). A specific closed-user group can be requested by supplying a facility code of 00000011 followed by the 8-bit closed-user group number.

User data. Data supplied by the calling DTE that is passed, unchanged, to the called DTE.

The Call Request packet arrives at the called DTE as an Incoming Call packet. If the DTE can accept the call, it returns a Call Accepted packet, which is received by the calling DTE as a Call Connected packet. The Call Accepted/Connected packet format is shown in Figure 10.24. The overall flow of information in the X.25 call establishment phase is presented in Figure 10.25.

At this point, data transfer may commence. However, if the call cannot be accepted, the DTE called responds with a Call Clearing packet (see below). Possible reasons for non-acceptance include the called DTE's refusal

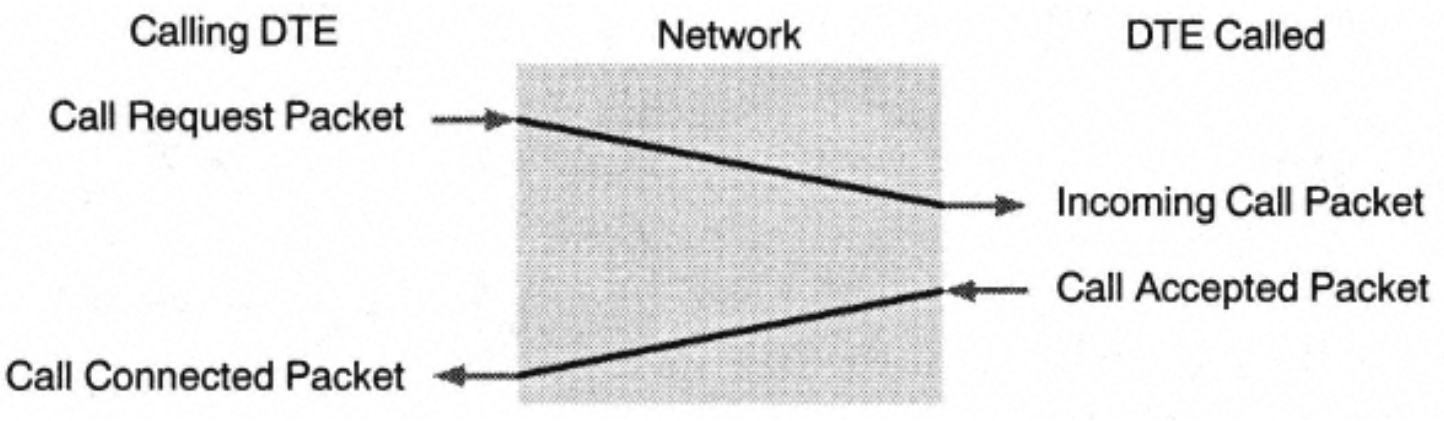

Figure 10.25
The steps in the call establishment phase.

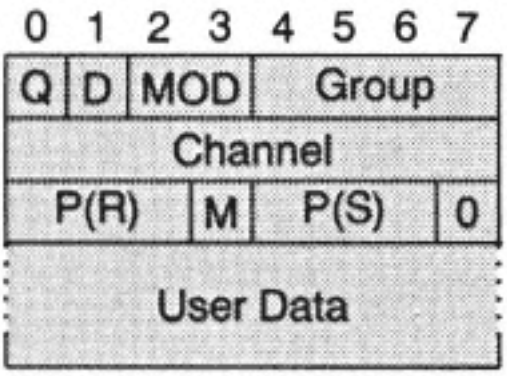

Figure 10.26
The X.25 Data packet.

to accept reverse charging or that the called DTE itself was in the process of making a call.

Once the virtual circuit has been established, the data transfer phase begins. Unless otherwise specified, the virtual circuit is full-duplex, meaning that the DTEs may transmit at any time as long as their transmit window is open. To ensure that packets are not lost and to permit flow control, X.25 implements a go-back-N sliding window protocol for level 3 Data packets. Each Data packet is sent with both its sequence number and the sequence number of the next expected packet. The format and contents of a Data packet are shown in Figure 10.26.

The Q-bit is the data qualifier bit; when cleared, it signals the DTE to forward the User Data to the application associated with the virtual circuit number. However, when set, the Q-bit allows the DTEs to exchange DTE-specific control data, and in these situations the User Data is not supplied to the application associated with the virtual circuit number. The D-bit, when set, requests an end-to-end acknowledgment for this packet *after* it has been successfully delivered to the application by the DTE (X.25 acknowledgments are discussed below).

P(R), the packet receive sequence number, and P(S), the packet sent sequence number, are the next expected and current packet sequence numbers, respectively. Sequence numbering is performed modulo-8 (i.e., 0 through 7) or modulo-128 (i.e., 0 through 127), depending upon the value of MOD. Modulo128 requires that P(R) and P(S) each occupy one octet. The sequencing allows the DTEs to perform end-to-end sequence number checking as well as flow control by withholding acknowledgments.

The User Data field holds a maximum of 128 octets unless otherwise specified by the network. Other valid maximum data lengths must be a power of 2 taken from the following list: 16, 32, 64, 256, 512, and 1024 octets. The M-bit is set if the current data packet is full and a subsequent packet contains data that is to be concatenated to the current data. For example, if an application sends a 132-byte message, two data packets are sent: the first a packet of 128 octets (with the M-bit set), followed by a second packet of 4 octets (with the M-bit cleared). The remote DTE then assembles the 132 bytes before forwarding them to the remote application.

There are three end-to-end acknowledgment packets that correspond to

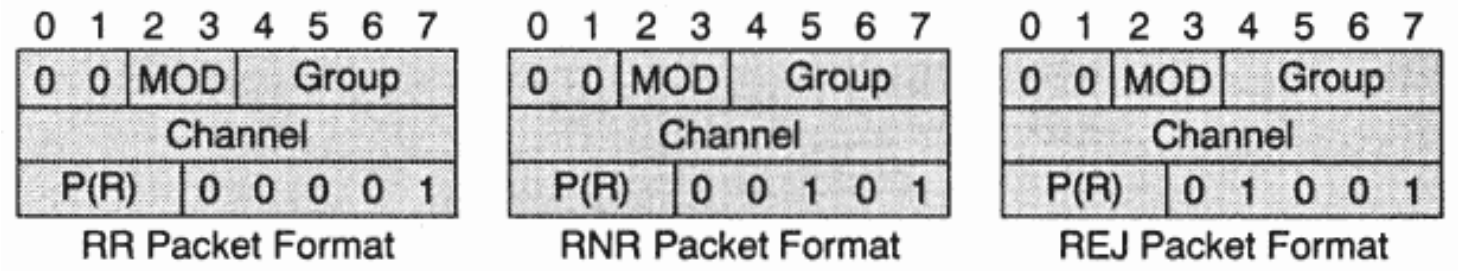
0 1 2 3 4 5 6 7
0 0 MOD Group
Channel
P(R) 0 0 0 0 1
RR Packet Format
0 1 2 3 4 5 6 7
0 0 MOD Group
Channel
P(R) 0 0 1 0 1
RNR Packet Format
0 1 2 3 4 5 6 7
0 0 MOD Group
Channel
P(R) 0 1 0 0 1
REJ Packet Format

Figure 10.27
X.25 end-to-end acknowledgment packet formats.

the SDLC/HDLC go-back-N protocol acknowledgment frames: RR, RNR, and REJ. The packet formats are shown in Figure 10.27 (P (R) denotes the sequence number of the next expected packet).

In addition to the sending and receiving of sequenced data packets, X.25 also allows for the transmission and reception of unsequenced Interrupt packets. Interrupt packets convey signaling information (such as a break signal) on a virtual circuit. The response from the DTE and DCE receiving the Interrupt packet is an Interrupt Confirmation packet, shown in Figure 10.28.

The exchange of data takes place through data packets across the virtual circuit. Acknowledgments can be piggybacked on Data packets or sent as RR packets; a typical cycle is shown in Figure 10.29.

The call clearing phase can be initialized by either DTE issuing a Clear Request packet. The DCE connected to the DTE issuing the Clear Request responds with Clear Confirmation, freeing the virtual circuit. When

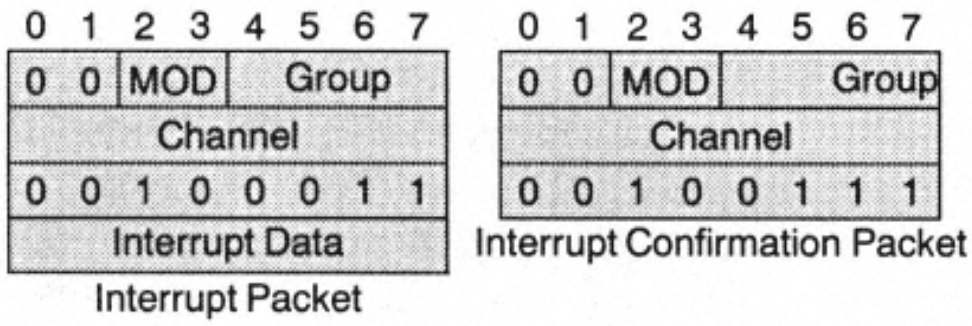
0 1 2 3 4 5 6 7
0 0 MOD Group
Channel
0 0 1 0 0 0 1 1
Interrupt Data
Interrupt Packet
0 1 2 3 4 5 6 7
0 0 MOD Group
Channel
0 0 1 0 0 1 1 1
Interrupt Confirmation Packet

Figure 10.28
The X.25 Interrupt and Interrupt Confirmation packets.

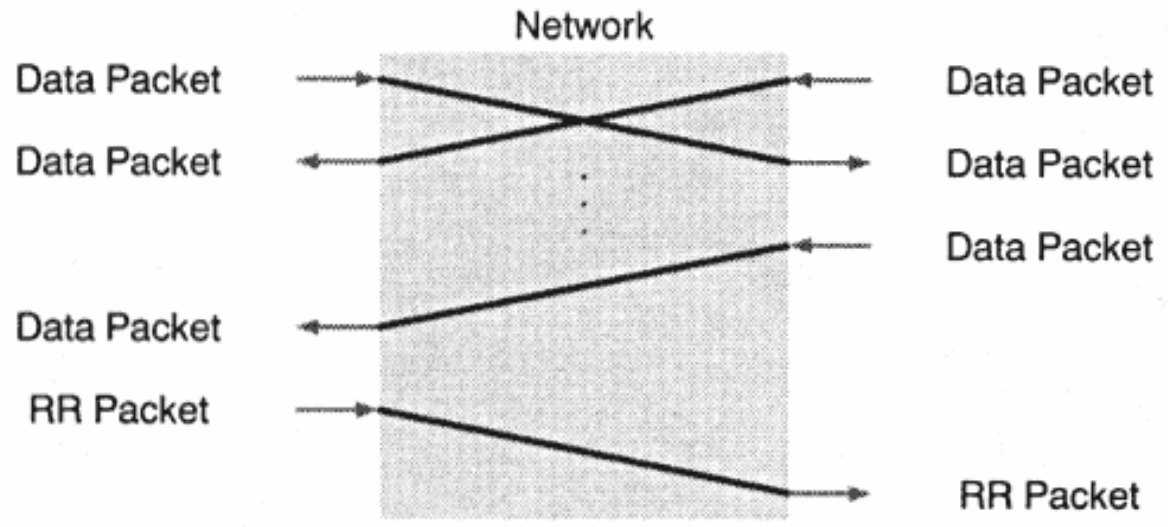

Figure 10.29
Data transfer using X.25.

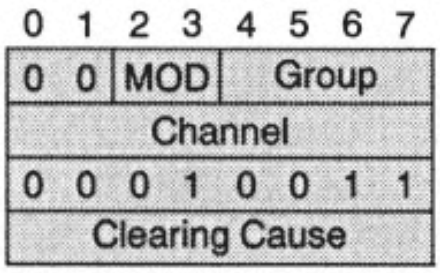

Figure 10.30
The X.25 Clear Request/
Clear Indication packets.

the packet arrives at the remote DCE, the DCE issues a Clear Indication packet to the remote DTE. The remote DTE then frees its virtual circuit and responds to its DCE with a Clear Confirmation.

The format of the Clear Request/Clear Indication packets is given in Figure 10.30, and the codes for the Clearing Cause field are given in Table 10.5. The format of the DTE and DCE Clear Confirmation packet is presented in Figure 10.31. Finally, the flow of packets for an X.25 call clearing is shown in Figure 10.32 (note that the Clear Confirmation packets are *local* rather than end to end).

In addition to the three phases described above, X.25 also defines two other phases, Reset and Restart.

The Reset phase allows a DTE to reset a specific virtual circuit: all sequence numbers are cleared and any data or interrupt packets associated with the circuit are discarded. Either DTE can reset a virtual circuit by issuing a Reset Request packet; the remote DTE resets the remote virtual circuit and responds with a Reset Confirm.

The Reset Request/Reset Indication packet format is presented in Figure 10.33. The codes for the Resetting Cause field are given in Table 10.6.

Table 10.5 X.25 Call Clearing Cause field values.

Code	0	1	2	3	4	5	6	7
DTE Clearing	0	0	0	0	0	0	0	0
Number Busy	0	0	0	0	0	0	0	1
Out of Order	0	0	0	0	1	0	0	1
Remote Procedure Error	0	0	0	1	0	0	0	1
Number Refuses Reverse Charging	0	0	0	1	1	0	0	1
Invalid Call	0	0	0	0	0	0	1	1
Access Barred	0	0	0	0	1	0	1	1
Local Procedure Error	0	0	0	1	0	0	1	1
Network Congestion	0	0	0	0	0	0	1	1
Not Obtainable	0	0	0	0	1	1	0	1

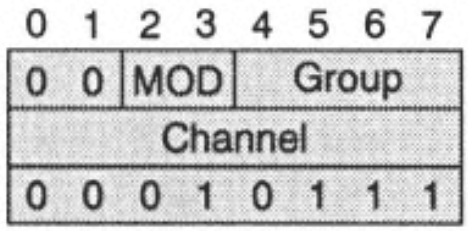

Figure 10.31
The X.25 Clear Confirmation packet.

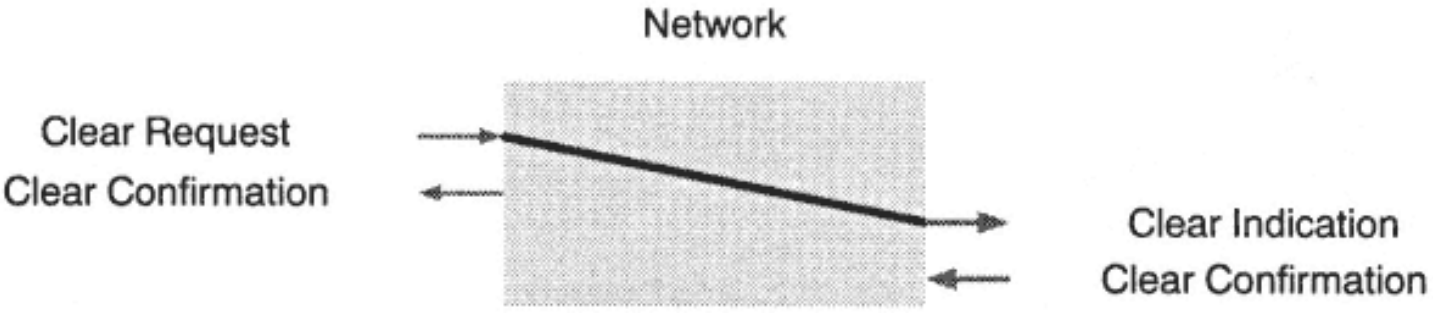

Figure 10.32
The steps in the Call Clearing phase.

The Diagnostic Code (Figure 10.33) has a value of 0. The format of the DTE and DCE Reset Confirmation packet is shown in Figure 10.34.

The Restart phase allows a DTE to clear *all* of its switched virtual circuits and to reset its permanent virtual circuits. Either the DTE or the DCE can request a restart by issuing a Restart Request packet; the response is a Restart Confirmation. Any switched virtual circuits must be re-established after a restart.

The format of the Restart Request/Restart Indication packet is as given in Figure 10.35. The codes for the Restarting Cause field in the Restart Request/Restart Indication packet are given in Table 10.7. The format of the DTE and DCE Restart Confirmation packet is shown in Figure 10.36.

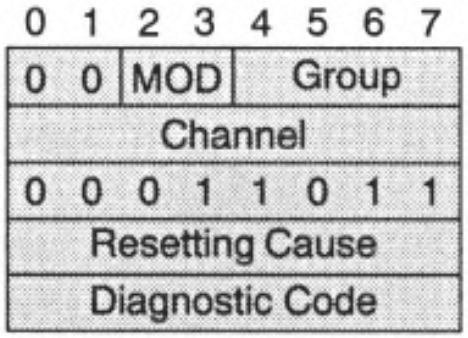

Figure 10.33
The X.25 Reset Request/
Reset Indication packets.

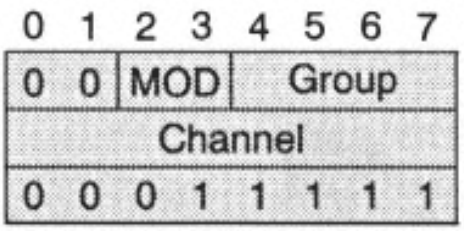

Figure 10.34
The X.25 Reset Confirmation packet.

Table 10.6 X.25 Resetting Cause field values.

Code	0	1	2	3	4	5	6	7
DTE Reset	0	0	0	0	0	0	0	0
Out of Order	0	0	0	0	0	0	0	1
Remote Procedure Error	0	0	0	0	0	0	1	1
Local Procedure Error	0	0	0	0	0	1	0	1
Network Congestion	0	0	0	0	0	1	1	1

Table 10.7 X.25 Restarting Cause codes.

Code	0	1	2	3	4	5	6	7
Local Procedure Error	0	0	0	0	0	0	0	1
Network Congestion	0	0	0	0	0	0	1	1

10.4.4
Network Access Protocols

From the above discussion on X.25, it is apparent that network access is handled by DTEs that support this protocol. The various X.25 levels can be embedded in a DTE's operating system, and users or applications requiring network access can follow the steps required by the specific operating system. Clearly, without some means of generating and recognizing X.25 packets, DTEs such as asynchronous terminals will be unable to access an X.25 network.

Fortunately, ITU-T has also developed a series of PAD protocols that permit non-X.25 DTEs access to wide area networks. These standards, X.3

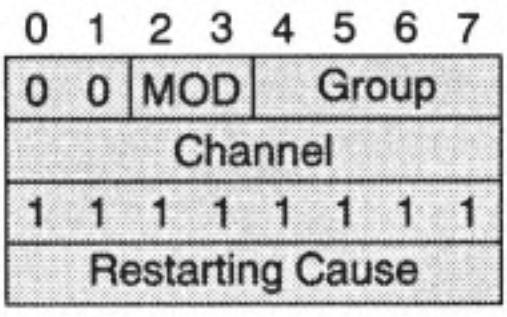

Figure 10.35
The X.25 Restart Request/
Restart Indication packets.

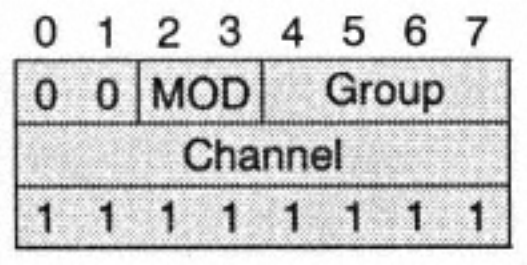

Figure 10.36
The X.25 Restart Confirmation packet.

(PAD parameters), X.28 (terminal-PAD interface), and X.29 (PADDTE interface) define how an asynchronous DTE (i. e., a terminal) can access a remote X.25 DTE through a packet switched network using a PAD.

Since different terminals and their users have different requirements, X.3 specifies a number of features that can be tailored to a specific terminal. (These features are maintained in a profile associated with the terminal for as long as it is connected to the network.) For example, the X.3 PAD parameters allow either the remote DTE or the local terminal to modify things such as

Whether the terminal can communicate with the PAD to change PAD parameters (PAD parameter 1).

Whether the PAD is responsible for echoing characters to the terminal (PAD parameter 2). If the terminal expects remote echoing but the DTE does not support it, the PAD can be requested to echo each character entered.

The character(s) used to signal the PAD that a packet is to be forwarded (PAD parameter 3). Packets can be forwarded when "full" (i.e., after 128 or 256 characters are entered), when a specific character is detected (such as a carriage return), or on each character (useful when working with a full-duplex screen editor).

Whether the PAD recognizes the flow control characters X-ON and X-OFF, allowing the user to suspend output to the terminal (PAD parameter 12).

If the PAD supports line-at-a-time input (that is, the end-of-line character causes the line to be forwarded in a packet). If so, the user should be allowed to make changes to the line before it is sent. X.3 defines editing character selection, in which the character delete, line delete, and line redraw keys can all be defined (PAD parameters 15, 16, 17, and 18).

A terminal connected to a PAD operates in one of two modes: communicating with the PAD (for example, for call setup or call clearing), or communicating with the remote DTE (for data transfer). When communication with the PAD is required (for example, to change one or more X.3 parameters), the user issues an escape sequence that is recognized by the PAD; data transfer then ceases while the user and the PAD communicate. The commands and associated responses are defined in X.28 and permit the user to establish a call, clear a call, and query and change PAD parameters. When data is to be sent to the DTE, the PAD is responsible for formatting it in an X.25 packet.

The DTE can also communicate with the PAD through the X.29 protocol, which defines a series of messages intended for use between the DTE and the PAD. These messages allow the DTE to query and select different PAD parameters as well as signal that the call is to be cleared. The ability to change PAD parameters is useful when, for example, the user activates a full-screen editor. At this time, the DTE can request that the PAD forward every character while the editor is in use.

The Q-bit in the X.25 packet allows the PAD to determine which packets are intended for the terminal (Q-bit is cleared) and which packets are intended for the PAD (Q-bit is set). X.29 packets are sent with the Q-bit set.

10.5 Examples of Wide Area Networks

Wide area networks fall into two broad categories.

The first category comprises those that have been developed by a company to work with that company's equipment. A good example of this is SNA, or Systems Network Architecture, developed by IBM, which consists of both hardware and software. SNA is a layered network architecture consisting of a physical layer (Physical Control); a data link layer (Data Link Control) that supports SDLC; and a network layer (Path Control). The Path Control Layer is responsible for packet routing and flow control.

DECNET is another example of hardware and software products developed by a manufacturer, in this case Digital Equipment Corporation, specifically for its own equipment. DECNET is part of Digital's DNA (Digital Network Architecture), a layered network architecture. It has a number of objectives, including the connection of Digital computers and operating systems, the support of any physical layer (i.e., full-duplex, half-duplex, synchronous, asynchronous, serial, parallel, and so on), and provision of facilities to support interprocess communications and file transfer. The Data Link Layer is supported by DDCMP (Digital Data Communication Message Protocol), a sliding window protocol with a window size of up to 127.

The second category comprises networks that conform to international standards, such as X.25. Most public packet switching networks, such as Datapac (Canada) and Transpac (France), support networking standards such as X.3, X.25, X.28, and X.29. There are a number of arguments for supporting international standards rather than the products of a single vendor. The most compelling is that supporting international standards permits the interconnection of equipment from different manufacturers. Another argument, to be pursued in later chapters, is that network interconnection can be easier if common protocols are adopted.

10.6 The Commkit Wide Area Network

Commkit is supplied with a layered model of a wide area network that demonstrates the functions of a Network Layer supporting virtual circuitsnotably, call establishment, data transfer, and call clearing. In addition, a sliding window protocol is implemented in the Data Link Layer. The wide

area network allows communication between PCs connected by null-modem RS-232-C cables. Each PC is treated as an individual node, and access to the network is through an application that allows the explicit testing of the Network Layer protocol.

The wide area network software consists of the following files:

`wanlow.c` the Physical and Data Link Layers. The Data Link Layer software supports a go-back-N protocol similar to SDLC/HDLC.

`wannet.c` the Network Layer, responsible for call establishment, data transfer, and call clearing.

`wanex.c` the application, an interactive message entry layer that allows the testing of virtual circuits.

`wanstrct.h` structures and state information specific to the Physical and Data Link Layer software found in `wanlow.c`.

`wandef.h` definitions and constants specific to the Network Layer, wannet.c.

`wancodes.h` definitions and constants used by both the Network Layer and the test software.

Internally, the software uses the `send()` and `recv()` primitives to support layering. The Data Link Layer, the Network Layer, and the application are all considered processes and hence are associated with unique queues.

10.6.1
Hardware Requirements

The Commkit wide area network software requires a minimum of two PCs, each with one (preferably two) serial ports and sufficient null-modem cables to allow the PCs to be linked in a bus-like fashion, as shown in Figure 10.37.

To facilitate routing, the wide area network software supplied with the Commkit diskette requires a specific ordering of port connections: serial port 1 must connect to serial port 2 (and vice versa); and a specific ordering of addresses: packets sent out port 1 are intended for PCs with lower-valued addresses, while packets sent out port 2 are intended for PCs with higher-valued addresses. For example, in Figure 10.37, PC D transmits to PC C through port 1, while PC C receives from PC D on port 2. The reasons for this are discussed below, and possible alternatives are left as exercises.

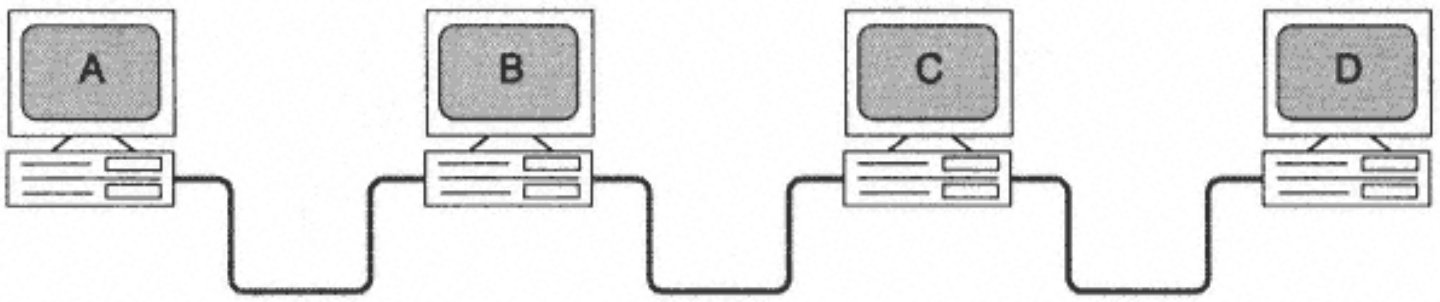

Figure 10.37
A 4-node Commkit wide area network.

10.6.2
Commkit Virtual Circuits

The Commkit wide area network consists of several interconnected PCs. Each PC acts as a network node while simultaneously permitting processes to communicate using virtual circuits. A Commkit virtual circuit is made from two distinct simplex paths through the network, connecting a process on either end.

To create one half of the virtual circuit, a process issues a Call Request packet addressed to the foreground process on a destination node. As the Call Request packet moves through the network, the Network Layer on each node creates a pointer to the entity that sent the packet (either a node or the original process). When the Call Request packet reaches the destination node, it is forwarded to the node's foreground process. At this point, a simplex reverse path has been established: from the process that received the Call Request packet, through the network layers on the intervening nodes, back to the process that originally sent the packet.

For example, in Figure 10.38, VC-1 is established by B sending a Call Request packet to A; similarly, VC-2 is established by A.

Since there can be a number of virtual circuits in operation at any time, each Network Layer maintains an array of them, `vcn_table`. An entry in `vcn_table` consists of three fields:

`dst_id`. This is the identifier of the queue associated with the source of the Call Request packet received by the Network Layer, one of `SP1IH, SP2IH`, and the calling process (typically `APPLICATION`). This identifier is stored, since it is this Network Layer's path *back* to the entity that transmitted the Call Request packet.

Data and Call Clearing packets put onto the queue indicated by `dst_id` are on the path back to the process that sent the original call packet.

`dst_vcn`. Each Call Request packet is sent with a virtual circuit number determined by the entity that sent it. If the entity is the calling process, the virtual circuit number is an identifier that allows the process to determine the source of the packet. However, if the entity that transmits the Call Request packet is a Network Layer, the virtual circuit number is the *subscript* of the Network Layer's `vcn_table` associated with this particular virtual circuit.

When the Network Layer receives a Call Request packet, the next avail-

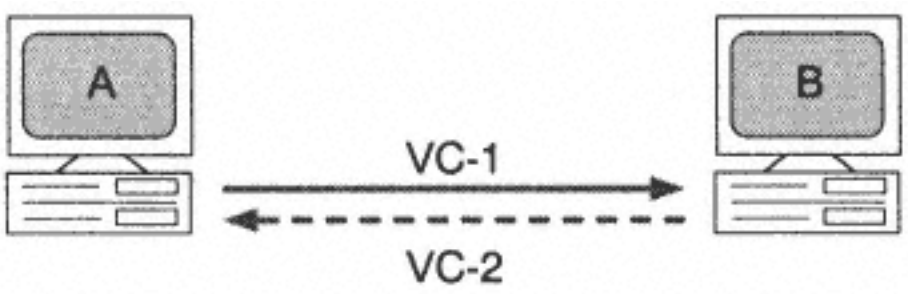

Figure 10.38
Two Commkit virtual circuits (each is a simplex channel).

able space in its `vcn_table` is found and the virtual circuit number is extracted from the packet and stored in the `dst_vcn` field. The value of the subscript is then placed in the Call Request packet as this particular Network Layer's virtual circuit number and sent to the next entity en route to the destination process.

When the network layer receives a Data or Call Clearing packet, the virtual circuit number is extracted from the packet and used as a subscript into `vcn_table`. The entry in `vcn_table` specifies the queue onto which the packet is to be placed (`dst_id`) and the packet's new virtual circuit number (`dst_vcn`). This cycle is repeated by each Network Layer.

`owner`. The destination entity of the Call Request packet as determined by the current Network Layer. The destination is one of `SP1IH, SP2IH`, and the called process (by default `APPLICATION`). It is determined by the Network Layer's routing algorithm.

The owner field is used in two situations. First, it ensures that the entity transmitting a Data packet is actually associated with this virtual circuit; if not, an Error packet is returned. Second, it verifies that the entity attempting to clear a call is allowed to do so.

For example, assume that process A on DTE X is to call process B on DTE Z and the route takes it through DTE Y. Initially, as shown in Figure 10.39, there is no path through the network from A to B.

Process A requests a path to process B by putting B's DTE address, Z, in a Call Request packet. The Call Request packet also contains a virtual circuit number generated by Ain this case, 0. The Call Request packet is then sent by A to its Network Layer.

Upon receipt of the Call Request packet, the Network Layer on DTE X determines that the route to Z is the link to Y (one of the serial ports). It then finds the first free entry in its virtual circuit table (in this example, the last, location 2) and in this location stores A as `dst_id, 0` as `dst_vcn`, and Y as owner. The virtual circuit number in the Call Request packet is changed to 2, the location of the entry in X's virtual circuit table. As shown in Figure 10.40, the first link in the virtual circuit has been established.

The Network Layer on DTE X then supplies the Call Request packet to the Data Link Layer for transmission to DTE Y. When the Network Layer on DTE Y receives the Call Request packet, the route is determined (out the

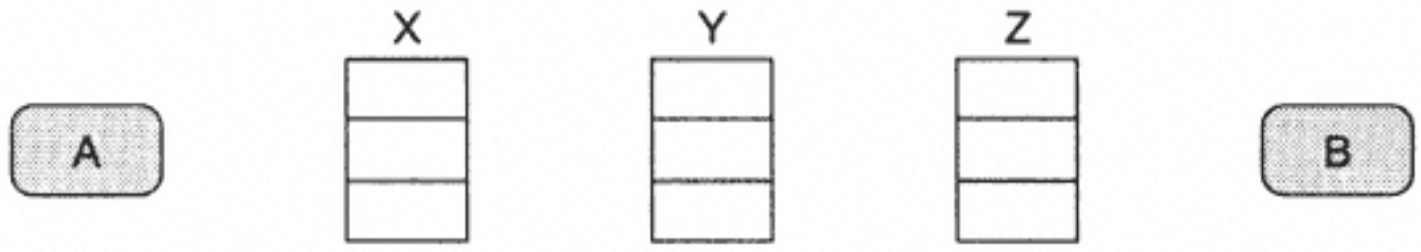

Figure 10.39
Initial network configuration.

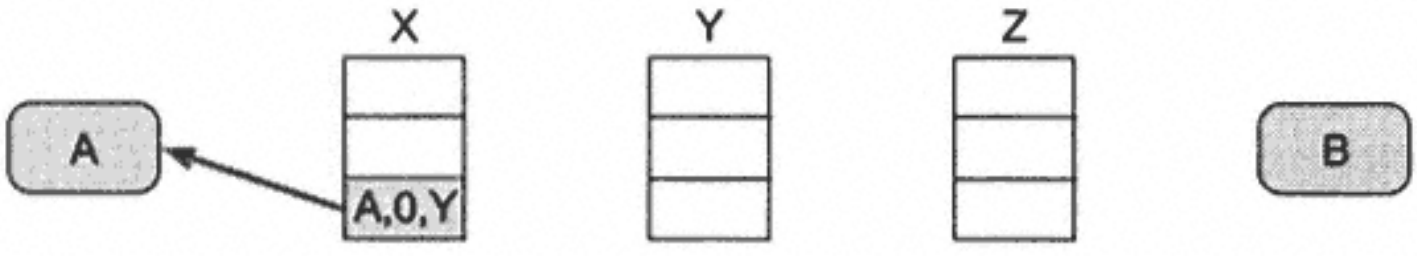

Figure 10.40
The first link in the virtual circuit: X to A.

link to Z), and the first free entry in DTE Y's virtual circuit table is found (in this example, location 0). The Network Layer stores X as `dst_id`, the virtual circuit number from the packet, 2 as `dst_vcn`, and Z as owner in location 0 of its virtual circuit table. The path, shown in Figure 10.41, now reaches to Y from A.

The virtual circuit number in the Call Request packet is changed to 0, and the packet is sent from DTE Y to DTE Z. Since the destination address in the Call Request packet is Z, the Call Request packet can be supplied to process B after the virtual circuit table is updated: `dst_id` is set to Y, `dst_vcn` is assigned 0, and owner is process B (in this example, location 1 is the next available entry in Z's virtual circuit table). The result of the changes is shown in Figure 10.42.

The Network Layer supplies the Call Request packet to process B with the virtual circuit number of 1. At this point a simplex circuit exists from process B to process A through the network; the completed virtual circuit is shown in Figure 10.43. Process B can now decide whether to accept or clear the call. If the call is accepted, process B issues a Call Request packet back to process A, creating a second path through the network.

Assuming that the call is accepted, if process B sends a message to process A, it must supply the virtual circuit number, 1, and the message in a Data packet to its Network Layer on DTE Z. The virtual circuit number, 1, is used as a subscript into Z's virtual circuit table; since the owner field specifies B,

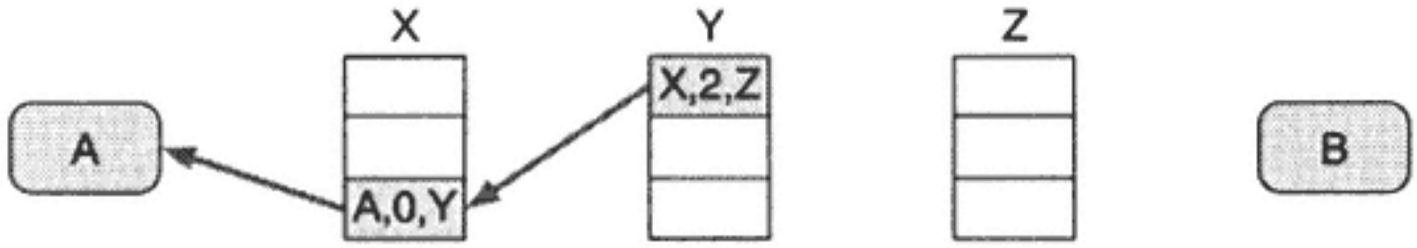

Figure 10.41
The second link in the virtual circuit: Y to X.

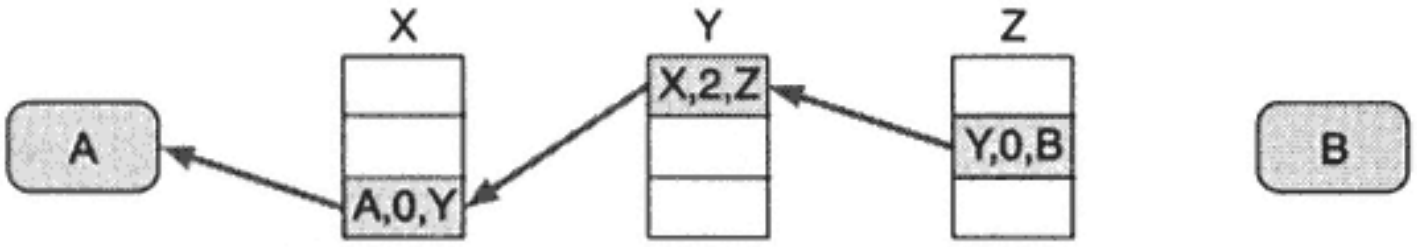

Figure 10.42
The third link in the virtual circuit: Z to Y.

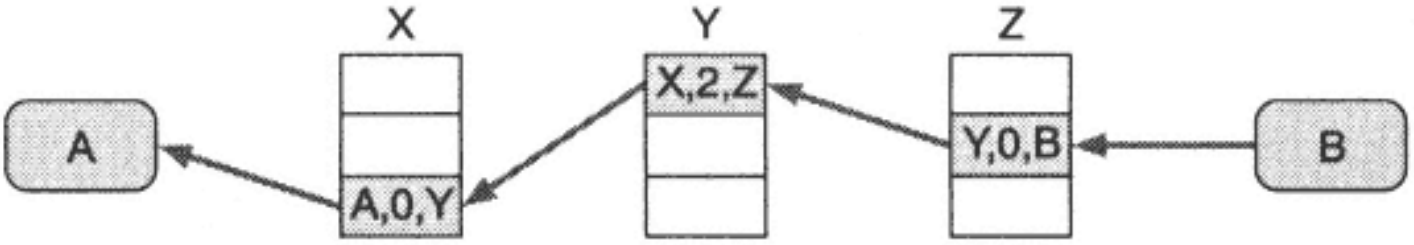

Figure 10.43
The completed simplex virtual circuit from B to A.

the transmission can take place. The virtual circuit number is changed to 0 and the Data packet is forwarded to the Data Link Layer that connects to Y.

When the Data packet arrives at Y's Network Layer, the virtual circuit number is extracted and once again used as a subscript. The Network Layer then checks the ownership of the entry and extracts the link, X, and the virtual circuit number, 2, from the virtual circuit table. After updating the packet's virtual circuit number to 2, Y's Network Layer forwards the packet to DTE X.

This cycle is repeated on DTE X, where the `dst_id` specifies the process A; the Network Layer on X forwards the Data packet to A with virtual circuit 0.

Call Clearing is initiated by the called process (in this example, B). The Call Clearing packet contains the virtual circuit number, thus allowing each Network Layer on the reverse path to remove the information from its virtual circuit table. For example, to clear the virtual circuit, process B sends a Call Clearing packet with virtual circuit number 1; Z's Network Layer ensures that B can clear this virtual circuit by checking the owner field associated with location 1 in `vcn_table`. The Network Layer proceeds to remove the virtual circuit information from the table and forwards the Call Clearing packet to Y with virtual circuit number 0. The cycle is repeated by each Network Layer all the way back to process A. Process A is then expected to issue a Call Clearing packet for its circuit to process B.

The Network Layer recognizes several error situations and handles them by returning an Error packet and an error code to the source of the packet. These error situations are

Invalid virtual circuit number. The source of the packet is not the same as the owner.

Unknown destination. The destination node does not exist.

Network congestion. A virtual circuit cannot be created because a virtual circuit table on one of the nodes is full.

10.6.3
Layers

The Commkit wide area network is implemented as a series of four distinct layers, shown in Figure 10.44.

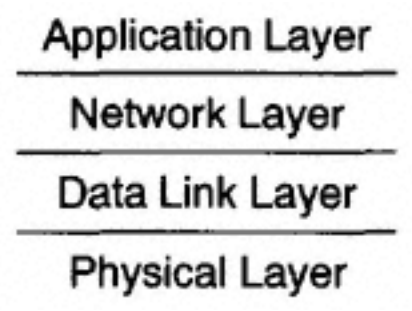

Figure 10.44
Commkit's wide area
network layering.

Application Layer

The Application Layer consists of a foreground process (`do_wan_ex()` in wanex. c) that allows the establishment and testing of virtual circuits. The process receives messages sent by either the Network Layer (`NETWORK`), the keyboard (`KEYIH`), or the clock (`CLKIH`) to the `APPLICATION` queue.

Keyboard characters are displayed on line 1 of the PC's screen and stored in the array buffer. The contents of the buffer are sent to the Network Layer when a carriage return is detected. The virtual circuit testing software recognizes five commands (the command is a single character stored in `buffer[0]`):

`R`send a call request frame; followed by a virtual circuit number, the destination node address, and the source node address.
`D`send a data transfer frame; followed by a virtual circuit number and the message.
`C`send a call clearing frame; followed by a virtual circuit number.
`T`enter test mode; followed by a virtual circuit number.
`S`stop test mode.

When test mode is entered, a data frame, `test_str`, containing the letters of the alphabet is sent by the test software to the specified virtual circuit. The virtual circuit must already be established. Packets are sent once a second if `x_count` is non-zero (`x_count` is the local process's transmit window used only in test mode). `X_count` is initialized to 5 when test mode is entered and is decremented each time a Data packet is sent, eventually closing the transmit window. The transmit window is reopened, by incrementing `x_count`, each time a packet is received.

The second byte entered after the command byte is the virtual circuit number (except in the case of S, which does not specify a virtual circuit). This number is entered as an ASCII character (assumed to be in the range 0 through 9); internally, this byte is stored in `buffer[1]` as the binary equivalent of the ASCII character (i.e., `0x00` through `0x09`). This allows the Network Layer to access its virtual circuit table without requiring special code for packets from the test process.

All messages are displayed in an output region, lines 5 through 23. Messages from the Network Layer are displayed in inverse video in the output region in the order they are received. Error codes are extracted from Error packets and displayed textually in the output region. As soon as a keyboard

message is sent, it is displayed in normal video in the output region on the next available line. `Update_line()` updates the line counter (`line`) and clears the next available line of the screen.

Network Layer

The Network Layer is a background process (`network_layer()`, found in `wannet.c`) associated with the queue `NETWORK` (a redefinition of `BACKGROUND_1`) that is responsible for the establishment and maintenance of virtual circuits.

`Network_layer()` is called when a packet is available on queue NETWORK. Packets are received from either of the serial ports (`SP1IH` or `SP2IH`) or from a process. The packet format shown in Figure 10.45 is the same in all cases (defined in `wandef.h`).

Four different packet codes are recognized: Call Request (`R`), Data (`D`), Call Clearing (`C`), and Error (`E`). The virtual circuit number, `vcn`, is a binary digit used as a subscript into the virtual circuit table, `vcn_table`. The contents of msg are ignored by the Network Layer except when the packet is a Call Request: `msg [P_DST]` and `msg[P_SRC]` denote the destination and source addresses, respectively. If the packet is an Error packet, `msg[ERR_CODE]` contains the error code.

Once the packet is obtained (using `recv()`), the code is inspected and one of the following functions is called (a packet with an unknown code is returned to its source as an Error packet with an error code of `UNKNOWN_CMD`):

`handle_call_request()`. Called whenever a Call Request packet is received (a code of `R`). It is responsible for obtaining a virtual circuit number and routing the Call Request packet to the next destination (either the link to the next DTE or the destination processby default `APPLICATION`). `Get_vcn_tbl()` returns the new virtual circuit number, which is stored in `vcn_tbl`. If there are no virtual circuits available, an error of `NETWORK_CONGESTION` is returned to the message source. The routing algorithm, when combined with the ordering of port connections described above, ensures that Call Request packets with unknown addresses will not circulate forever:

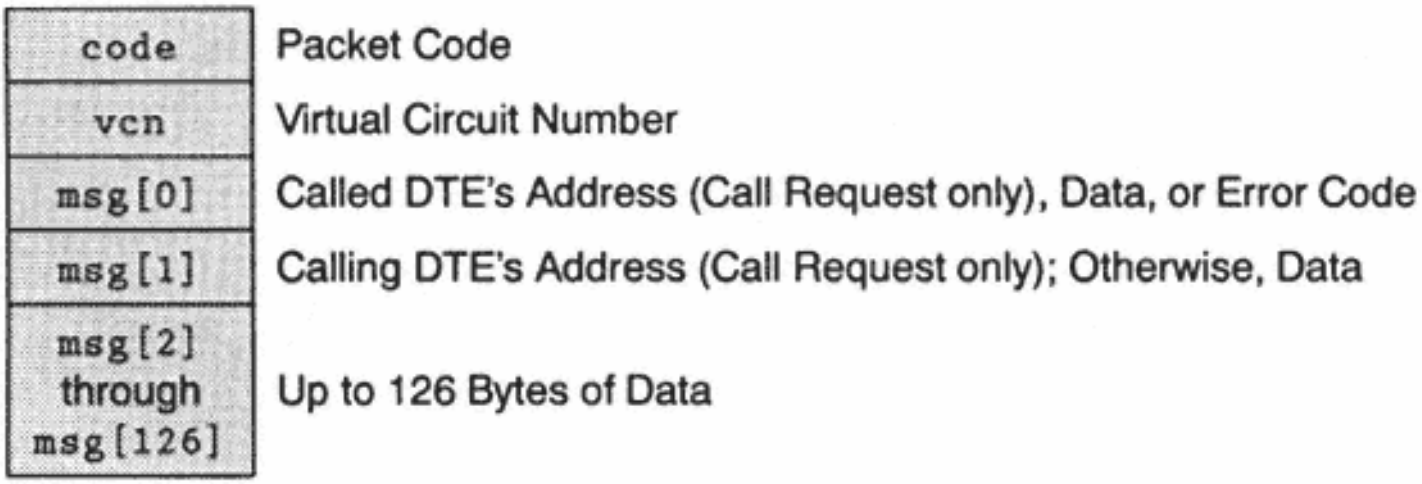

Figure 10.45
Commkit's Network Layer packet structure.

1. If the destination address is equal to the node's address, the destination of the packet is `APPLICATION`.

2. If the destination address is greater than the node's address and the packet is received from `SP1IH` or `APPLICATION`, the packet's destination is `SP2IH`.

3. If the destination address is less than the node's address and the message is received from `SP2IH` or `APPLICATION`, the packet's destination is `SP1IH`.

4. If the address does not exist, the Call Request is aborted and an Error packet is returned to the source of the packet, src, with an error of `UNKNOWN_DST`.

If the packet can be forwarded, the virtual circuit number in it is changed to `vcn_tbl` and the packet is sent to the destination determined by the routing algorithm.

`handle_data_transfer()`. Called by `network_layer()` when a packet with a code of `D` is received. The packet's `vcn` is used as a subscript into `vcn_table`; if the source of the message is the owner specified in `vcn_table`, the packet's `vcn` is changed to `dst_vcn` and the next destination is taken from `dst_id`.

An error code of `INVALID_VCN` is returned to the source of the message if the source is not the owner of this circuit.

`handle_call_clearing()`. Called when a Call Clearing packet is received (a code of C). The Call Clearing packet is forwarded to the next destination, and owner is set to `+1` to signal that this circuit is now free. An error code of `INVALID_VCN` is returned to the source if the source is not the owner of this circuit.

Data Link Layer

The Data Link Layer is responsible for the orderly transmission and reception of messages and acknowledgments across a link connecting two nodes. Reliable communications are achieved through the use of a sliding window protocol.

A single frame structure is supported by the Data Link Layer, consisting of a series of bytes enclosed by *STXETX*. The frame format is shown in Figure 10.46 (the size of each field in bytes is listed to the right of each field name). The individual fields of the frame are defined as follows.

The control field is a single byte indicating the type of frame. Two frame types are supported: supervisory and information. Frames are distinguished by the contents of the control field:

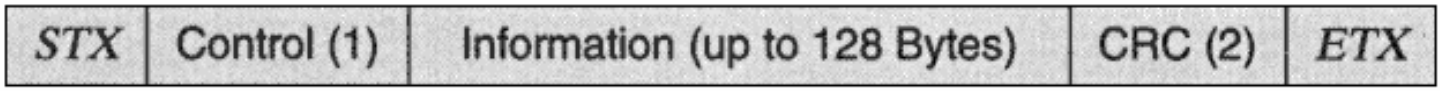

Figure 10.46
The Commkit Data Link Layer frame format.

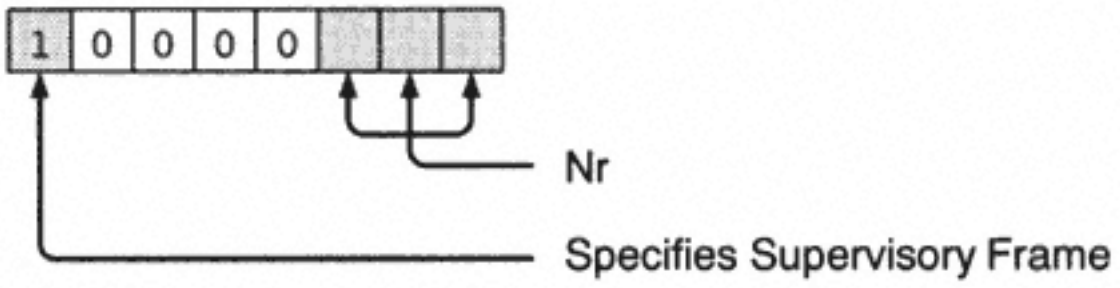

Figure 10.47
The Supervisory Control field format.

Supervisory. A supervisory frame is one that contains an acknowledgment (RR). The Information field is not sent in a Supervisory frame. The format of the Supervisory Control field is given in Figure 10.47.

Information. The Information frame Control byte indicates the sequence number of this packet (Ns) and the sequence number of the next expected packet (Nr). Information frames contain the Information field. Figure 10.48 shows the format of the Information frame Control field.

The Information field holds any message of up to 128 bytes in length. Transparency is supported by prefixing any *ETX* or *DLE* bytes with a *DLE* byte. The Information field is not sent in Supervisory frames.

Information frames are sent with a 3-bit sequence number (0 through 7), to which the receiving node responds with the *next* expected sequence number. Each frame sent is associated with a timer; if an acknowledgment is not received before the clock expires, the frame and all other pending frames are retransmitted. An Information frame transmitted is deemed received by the transmitting node when a Control field (either Supervisory or Information) is received with a sequence number greater than the frame's sequence number.

A 16-bit CRC, calculated using the cyclic redundancy check software in `crcsr.c`, includes the entire frame except for the leading *STX*. the trailing *ETX*. and any *DLE* prefix characters. It is sent with *all* frames and is calculated by the Physical Layer (see below).

The Data Link Layer is implemented as a series of procedures, the entry point being `data_link_layer ( ). Data_link_layer()` takes two arguments, `device` and `action`, and is activated by one of three conditions:

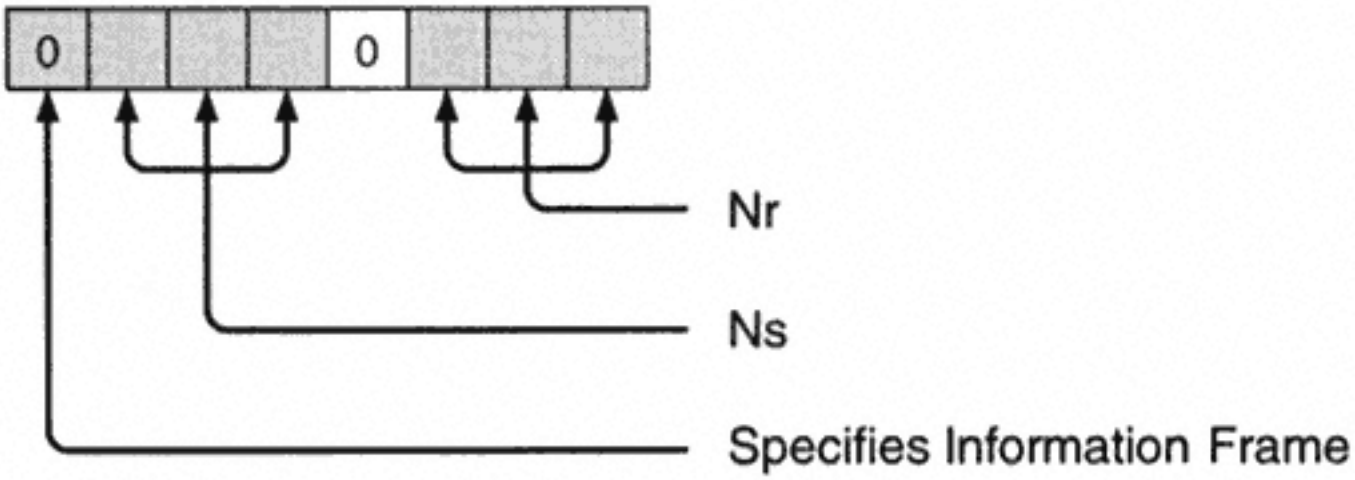

Figure 10.48
The Information frame Control field format.

A complete frame (either Supervisory or Information) being received by the Physical Layer (`device` is either `SP1IH` or `SP2IH`), indicated by the parameter `action` having a value of `MSGRECD`. Control is passed to procedure `d1_r_done()`.

The completion of a frame transmission (either Supervisory or Information) by the Physical Layer (`device` is either `SP1IH` or `SP2IH`), indicated by the value of action being `MSGSENT`. Control is passed to procedure `d1_x_done()`.

A 1-second clock pulse (from `low_level()`), indicated by the parameter device having a value of `CLKIH`. Control is passed to procedure `dl_clock()` . One-second clock messages are also sent from `low_level()` to the foreground process, `APPLICATION`.

The remaining three procedures used by the Data Link Layer are

`d1_clock()`. The clock procedure, called once a second to determine

If the time limit associated with the message currently awaiting an acknowledgment has been exceeded. When a message's time limit has been exceeded, that message and all other messages awaiting acknowledgment are retransmitted. The variable `wan.current_frame` is assigned the index of the message awaiting acknowledgment; transmission resumes from that message via a call to `dl_x_done()`.

If any messages supplied from the Network Layer are awaiting transmission. `D1_x_done()` is called if a message is on the device's queue (determined by examining the queue head of the device rather than issuing a `recv()` and the transmission window is open.

`dl_r_done()`. Called whenever a frame is received from either port.

Supervisory frames are assumed to contain acknowledgments; the acknowledgment sequence number (Nr) is extracted from the control byte, `wan.r_CTRL`. Procedure `update_ack()` determines the amount the window should be opened and updates `wan.awaiting_ack`.

Information frames contain both a sequence number and a piggy-backed acknowledgment. If the frame sequence number is equal to `wan.expected`, the message part of the frame is sent to the Network Layer and the acknowledgments are updated. The flag `wan.pending_ack` is assigned `TRUE` to ensure that if no Information frames are to be transmitted, at least a Supervisory frame will be sent to maintain throughput.

Pending frames awaiting acknowledgment are *not* retransmitted simply because their acknowledgment has not arrived. Instead, the `wan.current_frame` index is updated only if the frame awaiting an acknowledgment has been waiting for at least one second.

`dl_x_done()`. Called whenever a frame has been sent *or* a condition arises in which a frame can be sent. Frames (Information or Supervisory) are transmitted only if one of the following conditions holds (in order of priority):

1. Messages are in the process of being retransmitted.

2. A new message is available for transmission and the window is open.

3. A pending acknowledgment is awaiting transmission.

Since entry to these procedures occurs when interrupts are disabled, mutual exclusion is guaranteed.

Physical Layer

The Physical Layer is a device driver and is responsible for the transmission and reception of frames. It recognizes the Data Link Layer frame format and is responsible for calculating the CRC associated with each frame. The entry point to the Physical Layer is the procedure `physical_layer()`, which is called by `low_level()` and takes three parameters:

`port`. The port (either `SP1IH` or `SP2IH`) associated with the `action`.

`action`. A code indicating the result of the function just completed by the serial port, either `XMITDONE` or `RECVDONE`.

`in_ch`. The character just received (if `action` is `RECVDONE`); otherwise, 0.

Each port is associated with its own data structure in the array `wan` (`wan[0]`

for `SP1IH` and `wan[1]` for `SP2IH`, defined in `wanstrct.h`).

The Physical Layer is full-duplex and can be sending and receiving out both ports at any particular moment. It can be described in terms of its two primary functions, the transmission and reception of frames.

A transmission (`action == XMITDONE`) is initiated by the Data Link Layer sending an *STX* out one of the ports and supplying the Physical Layer with the following information (note that `wan` refers to either `wan[0]` or `wan[1]`, depending upon the value of port):

1. The frame's control byte, `wan.x_CTRL`, indicating whether the frame is Supervisory or Information. The control byte is included in the CRC.

2. The size of the message to be transmitted, `wan.x_size`. A size of 0 indicates that there is no message (i.e., this is a Supervisory frame). Any other value is taken to mean that there is a message to be sent. `Wan.x_count`, initially0, is incremented as each byte is sent and compared to `wan.x_size`.

3. A pointer to the message, `wan.msgptr`.

The transmission part of the Physical Layer is implemented as a state machine consisting of eight states (the current state is indicated by `wan.x_state`):

`SEND_CTRL`. This state is entered after transmission of the *STX* has been signaled. The control byte, `wan.x_ctrl`, is sent and included in the CRC. The next state depends upon the message size; if wan. `x_size` is 0, `wan.x_state` changes to `SEND_CRC1`; otherwise, to `SEND_DATA`.

`SEND_DATA`. The next byte in the message (pointed to by `wan.msgptr`) is to be sent and included in the CRC. The byte count, `wan.x_count`, is incremented as each byte is sent. When the count is equal to the

wan.x_size, the state changes to SEND_CRC1. If the byte's value is either *DLE* or *ETX*. the byte is saved in wan. next_ch, the next state (either SEND_DATA or SEND_CRC1) is saved in wan.next_x_state, and the state changes to SEND_NEXT.

SEND_NEXT. An inserted *DLE* has just been sent. The byte to be sent in this state is taken from wan. next_ch and the value of the state changes to the value in wan.next_x_state.

SEND_CRC1. The lower eight bits of the CRC are to be sent. If the CRC byte is neither *DLE* nor *ETX*. the next state is set to SEND_CRC2. However, if the byte's value is either *DLE* or *ETX*. the byte is saved in wan. next_ch, SEND_CRC2 is saved in wan.next_x_state, and the state changes to SEND_NEXT.

SEND_CRC2. This state is identical to SEND_CRC1, with the exception that the upper eight bits of the CRC are sent. The next state is SEND_ETX unless the byte has a value of *DLE* or *ETX*. in which case the state changes to SEND_NEXT.

SEND_ETX. The frame-ending *ETX* is sent. The state changes to SEND_DONE.

SEND_DONE. This state is entered after the final *ETX* has been sent, indicated by the UART signaling XMITDONE after the *ETX* is sent. The Data Link Layer is called to signal that the transmission has completed. The state then changes to SEND_IDLE.

SEND_IDLE. This state indicates that the transmitter is idle.

DLE insertion occurs whenever an *ETX* or *DLE* is found in the message *or* in the pair of CRC bytes. The inserted *DLE* is not included in the CRC; however, the byte following the *DLE* is part of the CRC.

The receive state machine (action == RECVDONE) is entered whenever data is received from either of the serial ports. Initially, the state machine is searching for an *STX*. denoting the start of frame. Once the frame-ending *ETX* is found, the message part of the frame and the control byte are returned to the Data Link Layer for processing.

The different receive states (indicated by wan.r_state) are as follows:

WAIT_STX. This is the reception idle state, entered initially and thereafter whenever a frame has been received. If an *STX* is found, both the receive CRC (wan.r_crc) and the receive data count (wan.r_count) are cleared. The state changes to WAIT_CTRL.

WAIT_CTRL. The byte following the *STX* is assumed to be the control byte; it is stored in wan. r_CTRL and included in the receive CRC. The state changes to WAIT_DATA.

WAIT_DATA. In WAIT_DATA, all incoming bytes except *ETX* and *DLE* are included in the incoming CRC and stored in the receive message buffer (wan.r_msg). The detection of a *DLE* causes a state change to WAIT_NEXT; the *DLE* is not included in the CRC.

When the frame-ending *ETX* is found, the number of bytes in the message is decreased by two (since the last two bytes are the CRC). The calculated value of the CRC should be 0 if no errors were detected. The reception of an error-free packet causes control to pass to the Data Link Layer; otherwise, the packet is discarded. The receive state then returns to `WAIT_STX`.

`WAIT_NEXT`. This state is entered after a *DLE* is detected in the `WAIT_DATA` state. The byte received in this state is included in the CRC and stored as part of the message. Control returns to `WAIT_DATA`.

If, for some reason, the frame-ending *ETX* is lost, the receive state machine will stay in either the `WAIT_DATA` or `WAIT_NEXT` state. Subsequent frames will be included in the incoming message buffer and hence lost. More seriously, a subscripting error will occur if the incoming data is simply copied into the message buffer. To avoid this situation, before a byte is written to the message buffer the current receive count is checked; if the value exceeds the maximum message size, the receive state is returned to `WAIT_STX`.

10.7
Using the Commkit Wide Area Network

The Commkit wide area network software is supplied on the Commkit diskette, both as source listings and as the executable, `wanex.exe`.

10.7.1
Creating the Wide Area Network Software

The Commkit diskette is supplied with both source and executable versions of the wide area network software that allows keyboard data entry. If any of the wide area network modules (`wanex.c`,`wannet.c`, or `wanlow.c`) are changed, the modules should be recompiled to make a new executable, `wanex.exe`.

The executable is created using the make utility:

```
C:\> make wanex.exe
```

10.7.2
Running the Wide Area Network

Once an executable is obtained (either from the diskette or by recompiling), it can be run by typing `wanex` after the MS-DOS prompt and specifying the line speed, the address of the node, and the transmit window size used by the Data Link Layer. For example, to set the PC up as node S on a 1200-bps network with a window size of 3, type

```
C:\> wanex 1200 S 3
```

At this point, the screen is cleared and communications can proceed. Data entered by the user is displayed on line 1 of the screen; all or part of a line can be erased using the backspace key. When the carriage return is pressed, line 1 is erased, but the line is displayed on the next available line in the output region (lines 5 and 23) in normal video. As well as displaying keyboard input, the output region displays all messages and diagnostics from the Network Layer in inverse video. The output region "wraps around" when the next line to be displayed exceeds its last line. Control can be returned to MS-DOS at any time using CTRL-C or CTRL-Alt-DEL.

If several PCs are to be connected to form a wide area network, the cables should be connected from port 1 to port 2, as shown in Figure 10.49. Remember that the address assigned to each station depends upon its position with respect to other stations.

10.7.3
Testing Virtual Circuits

`Wanex.exe` allows the testing of virtual circuits across the Commkit wide area network. Three virtual circuit phases and related packet structures are supported.

Call Setup

A process on one node can attempt to establish a virtual circuit with a second process on a remote node. Calls can be established at any time (for example, during the data transfer phase) as long as there are virtual circuits available. Data transfer can take place only after a Call Request packet is received.

Half of a virtual circuit is established by a user explicitly creating a Call Request packet, the format of which is shown in Figure 10.50. The packet's fields are defined as follows:

R (uppercase *R*). The code used to indicate a Call Request packet.

VCN. The virtual circuit number assigned by the user for this particular virtual circuit. It can be any keyboard character. The virtual circuit number assigned by the user allows the user to distinguish between incoming

Figure 10.49
Ordering of ports in the Commkit wide area network.

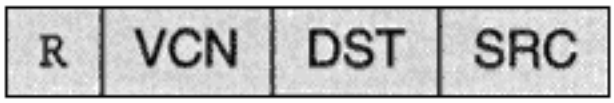

Figure 10.50
Commkit's Call Request packet.

calls. The VCN is not transmitted across the network; however, all incoming packets for the specific circuit have the original VCN displayed as part of the message.

DST. The address of the destination PC, a single character. It need not be of an existing node (see below).

SRC. The address of the source node, a single character. This should be the address of this station, to allow the remote to respond.

A call request packet establishes one half of a virtual circuit (i.e., from the DST back to the SRC); the node called must respond with another Call Request packet to the calling node if a full-duplex virtual circuit is to be established.

As an example, consider a wide area network consisting of two PCs, one with address S, the other with address E. For a virtual circuit to be established between these two nodes, one (in this example, S) issues a Call Request packet. The virtual circuit number chosen is the letter *a*. meaning that all subsequent packets *received* on this virtual circuit will contain *a* as the virtual circuit number. The user on node *S* therefore types *RaES* as the Call Request packet (*R* indicates a Call Request packet; *a* is the virtual circuit number identifying this circuit; *E* is the address of the destination; and S is the address of the source). When carriage return is typed, line 1 clears and *RaES* appears in the output region.

After a moment, a Call Request packet appears on node E in inverse video on the output region as *R0ES*. This is interpreted as a Call Request packet from the process on node S with the virtual circuit number 0; all packets sent with 0 will appear on node S with the virtual circuit number replaced by *a*. This is virtual circuit number 0 on node E because it is the first circuit established.

To make a full-duplex circuit, the user on node E must respond with a Call Request packet. Assuming that the virtual circuit number chosen by the user on node E is the letter *x*, the Call Request packet must be entered as *RxSE* (i.e., a Call Request packet on virtual circuit *x* to node *S* from node E). Any packets on this half of the virtual circuit will contain the virtual circuit number *x*.

This Call Request packet arrives at node S and is displayed in inverse video as *R*1*SE*. The virtual circuit number appears as 1 because circuit number 0 was used for the outbound call and 1 is the next available circuit number on node S.

A full-duplex virtual circuit has now been established between node S (transmitting on VCN 1 and receiving on VCN a) and node E (transmitting on VCN 0 and receiving on VCN x). The pair of (simplex) virtual circuits are represented diagrammatically in Figure 10.51.

If the destination address supplied does not exist between the two nodes, an error of `UnknownDST` is displayed in inverse video. If the virtual circuit

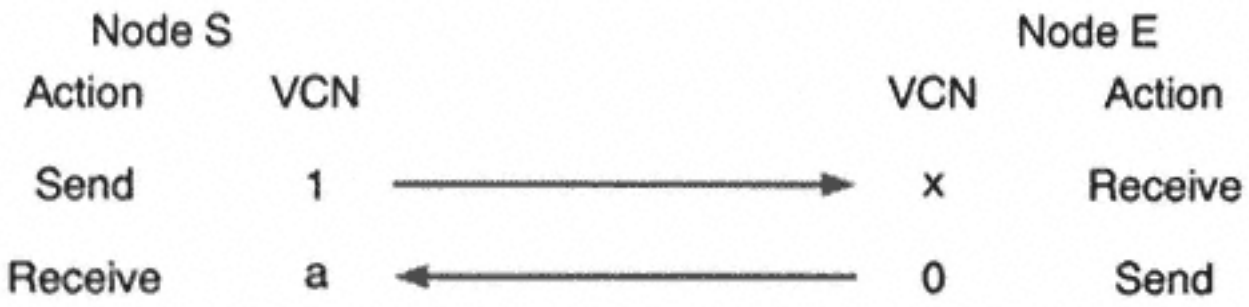

Figure 10.51
A pair of simplex virtual circuits between nodes E and S.

table is full on one of the nodes, the call request will be blocked and an error message `Network Congestion` is displayed, also in inverse video.

Data Transfer

Once a call has been established and a virtual circuit number returned, data transfer can take place between the two processes at either end of the circuit. A Data packet consists of the character *D* (denoting a Data packet), a virtual circuit number (a character), and a string of characters making up the message, terminated by a carriage return. Figure 10.52 is an example.

Data is displayed in the output area of the destination node in inverse video. The virtual circuit number is the value specified by the user when the Call Request is established.

Continuing with the example begun in the Call Request phase, if the user on node E is to send a message to the user on node S, the message must be sent in a packet with virtual circuit number 0. The message arrives at node S and is displayed in inverse video with the virtual circuit number specified by the user; in this example, the message would appear with virtual circuit number *a*.

For example, to send the message "Hot this summer, isn't it?" from node S to node E, a Data packet must be constructed, consisting of the data transfer character D, the virtual circuit number (in this case, 1), and the message. This is entered on line 1:

```
D1Hot this summer, isn't it?
```

After a moment, the packet appears in the output area of node E in inverse video as

```
DxHot this summer, isn't it?
```

where x is the virtual circuit number on node E.

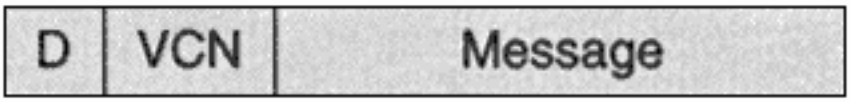

Figure 10.52
A Commkit Data packet.

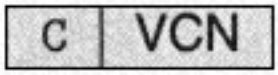

Figure 10.53
The Commkit Call
Clearing packet.

If the virtual circuit number in the Data packet is not associated with the process, an error message of `Invalid VCN` in inverse video, followed by the virtual circuit number, appears in the output area.

Once a virtual circuit has been established, the test software can be made to send a stream of Data packets to the process at the other end of the virtual circuit. This is known as test mode and is initiated by typing *T* followed by the virtual circuit number over which the transmissions are to occur. Each packet contains the lowercase letters of the alphabet (*a* through *z*).

To avoid flooding the node's queues, the test mode packets have a window size of 5; if a Data packet is not received after five packets have been sent, the window closes. Transmissions resume as Data packets are received, opening the window. Both nodes can be in test mode simultaneously.

For example, if node E is to enter test mode and send packets across the virtual circuit, the user types *T0*. After a moment, the Data packets will start to appear on node S.

Test mode is disabled by typing *S* followed by a carriage return.

Call Clearing

A call can be initiated by either end of the virtual circuit issuing a Call Clearing packet along with the number of the virtual circuit to be cleared. The format of the Call Clearing packet is shown in Figure 10.53.

The virtual circuit number to be used is the one in which all Data packets have been sent. In the example used above, if the user on node S is to clear the call, the virtual circuit number to be specified is 1. This appears on node E's screen as `Cx` in inverse video, indicating that no further data will be received from virtual circuit x.

Upon receipt of a Call Clearing packet, the user should respond with a Call Clearing packet, thereby freeing up the virtual circuits making up the other half of its virtual circuit.

If the virtual circuit number is not in use by this process, the request is aborted by the Network Layer and a diagnostic of `Invalid VCN` in inverse video, followed by the virtual circuit number, is displayed.

10.7.4
Low-Level Testing

The virtual circuits requested by the user and established by the Network Layer are supported by the Data Link and Physical Layers. Communications between each pair of nodes is supported by the sliding window protocol

Table 10.8 Supervisory frame field format.

Nr	Internal Value	Displayed as
0	10000000	C
1	10000001	u
2	10000010	e
3	10000011	â
4	10000100	a
5	10000101	a
6	10000110	a
7	10000111	c

discussed in Section 10.6.3. This protocol is hidden by the Network Layer, so to examine the actions associated with the protocol it is necessary to use the line analyzer.

By placing the line analyzer between two PCs running the wide area network software, the different Data Link frame structures can be seen. All frames have the same format, an *STX* followed by a number of bytes (two of which are the CRC) and finally an *ETX*.

The two frame types are

Supervisory. Supervisory frames are short, consisting of a single control field containing the value of Nr, and will appear on the line analyzer as one of the bytes shown in Table 10.8.

Information. Any string of bytes sent across the channel is an Information frame. The second byte (after the *STX*) is the Control field and contains the values of Ns and Nr. The line analyzer permits the decoding of the Control field. For example, a ♣ (0x05) indicates an Nr value of 5 and an Ns of 0.

The packet's code byte is the third byte to be transmitted (after the *STX* and Control field) and has one of the following values: R (Call Request), D (Data), C (Call Clearing), and E (Error). The remaining bytes in the frame depend upon the type of packet. The two bytes prior to the *ETX* are the CRC (more bytes may appear if the CRC contains a *DLE* or *ETX*).

The piggybacking of acknowledgments is best illustrated by putting both nodes into test mode. Once the screen is filled, stop the output using CTRL-S.

If a station has been changed from a line analyzer to a node (or vice versa), it is necessary to restart all the machines on the network to ensure that the nodes exhibit the correct values of Ns and Nr. Failure to observe this can result in one or more nodes terminating without restoring DOS interrupts (time to hit the Reset button).

10.8 Summary

A wide area network (WAN) is a collection of nodes interconnected by a series of links. There is no single WAN topology, although most are partially connected meshes. The purpose behind a wide area network is to offer users a shared communication facility that is made available on demand. This is different from the telephone system, which must devote an entire circuit to a pair of DTEs for the duration of their communication.

Wide area networks achieve shared communication through packet switching, a technique in which information is transmitted in a data structure known as a packet that is routed through the network from node to node. Routing can be achieved through any number of techniques, but most commercial WANs use some form of routing table.

There are three types of packet service that a wide area network can offer:

Datagrams, in which the packet is put on the network and the various nodes make a best effort to have the packet reach the intended destination. There is no guarantee that the datagrams will reach their intended destination or that they will reach it in the order they are sent.

Messaging, in which a single message is broken into a series of packets and transmitted between nodes. Unlike a datagram, a message is reassembled at each node at which it arrives; the node is responsible for determining the next leg in the route and scheduling a time for the transmission of the message.

Virtual circuits, which closely resemble a telephone call, consisting of a number of distinct steps. In the first step, call establishment, a Call Request packet moves from node to node, defining a path (the virtual circuit) through the network. Once the path is established and both ends agree to communicate, the second step can be instituted: the transfer of data across the route defined by the Call Request packet. When either or both DTEs agree to terminate the call, a Call Clearing packet can be sent, freeing the different stages across the network.

Wide area networks illustrate two important concepts in data communications: layering, in which software is divided into discrete parts, each performing a specific task; and standards, which ensure that DTEs from different manufacturers can communicate. A good example of a layered wide area network standard is ITU-T's X.25, a three-layer architecture defining a Physical Layer, a Data Link Layer, and a Network Layer.

10.9 Exercises

1. (A) Why isn't it necessary for the Data Link Layer protocol to be reliable in a datagram network?

2. (A) How does X.25 distinguish Data packets from other packets?

3. (A) Show the layers necessary for a messaging system and for a virtual circuit.

4. (A) The description of how a simplex channel is established on Commkit showed one half of the connection (from A to B). Complete the virtual circuit by making the connection from B to A and showing the virtual circuit tables in nodes X, Y, and Z.

5. (B) Connect two PCs running the wide area network software via a third running the analyzer. Send nine messages from one station to the other, recording the responses. Explain why the Control field in the Supervisory frames appears as follows:

ü é â ä à á ç Ç ü

The frame delimiters are displayed on the line analyzer as ☺ (*STX*)and ♡ (*ETX*). The CRC bytes are within the frame (i.e., before the final *ETX*).

6. (B) The Control field of an Information frame can also be represented using the PC's character set. Using the same hardware configuration as above, complete the table shown in Table 10.9 for the Control fields of an Information frame.

7. (B) Show by means of a diagram that the routing algorithm used by the Commkit wide area network actually works. Specifically, set up a network of two stations, A and C, and show what happens when A attempts to send a message to B. Repeat the exercise for C sending a message to B.

8. (B) Connect two PCs, one running the line analyzer and the other running the wide area network. Send a message from the wide area network PC, making sure that the cables are connected so that the message appears on the line analyzer. What is the maximum number of messages that can be sent? Explain the limit; how can it be changed? Why are the messages sent in order?

9. (A) Draw the transmission and reception state diagrams for each layer of Commkit's wide area network. Compare these diagrams to those describing the stop-and-wait protocol in Chapter 6.

Table 10.9 Table for Exercise 6.

	Ns							
Nr		≻		0	@		'	p
			!	1		Q	a	
			$					
	♣	§	%					
	♠							
								w

10. (C) The previous exercise highlights a limitation of the wide area network software: it doesn't know when to give up. Modify the software so that if a response it not detected after ten retransmissions of the same message, the remote node is assumed to be dead and the communication is abandoned.

11. (C) The wide area network supplied with Commkit requires that the stations be arranged in an order that ensures that all stations "downstream" of a station's secondary port have addresses greater than that of the station, while stations "upstream" of a station's primary port have addresses less than those of the station. This will clearly lead to a number of annoying problems should someone try to add a new station between two existing stations with consecutive station numbers. For example, how can a third station be added between a station with address A and a station with address B? The answer is, not surprisingly, it can'tnot without reordering the addresses associated with one or more existing stations. It is not advisable to change the addresses associated with stations within any type of network (including wide area), since an address may be "known" by any number of other stations. Suggest a reason why this is so.

12. (C) The protocol used by the Data Link Layer is a partial implementation of the SDLC protocol. Information frames are supported, and so are receive ready (RR) frames; however, reject (REJ) and receive not ready (RNR) are not. Similarly, if a packet is received correctly but is out of sequence, the software does not respond; rather, it waits for the transmitter to send the entire sequence of messages again.

Modify the protocol to support the following:

> REJ Control frames, which are typically sent when a packet has been received with a CRC error. The value of Ns indicates the next valid frame number.
>
> RNR frames, which should be sent when the receive cannot accept packets.
>
> An RR frame if an Information frame has been received correctly but out of sequence. The value of Ns indicates the next valid frame number.

13. (C) The Physical Layer is written assuming that once a reception starts, there will always be an *ETX* somewhere in the incoming data stream, allowing the reception to conclude. If an *ETX* is lost, it is also assumed that another frame will be sent, causing the message buffer to reach its limit, forcing the reception state machine back to the WAIT_STX state. However, if the transmitting station should cease transmission (i.e., someone turns it off or the cable is pulled), the reception software will hang, waiting for an *ETX* that will never appear.

Redesign the reception part of the Physical Layer to ensure that if a byte is not received within five seconds, the reception software is to "assume"

that the transmitter is no longer transmitting. A message code should be returned by the Physical Layer to the Data Link Layer, the Network Layer, and finally the Process Layer, informing the user of the problem. Once implemented, how can the changes be tested? Does the software support temporary loss of transmission (for example, a loss of less than five seconds)?

14. (C) The previous exercise has other ramifications: how does the Data Link Layer inform the Network Layer that an error has occurred? Once the Network Layer has been informed, it should cease all communications out that link. Modify the software to support this feature.

This fixes another problem with the software: Calls to nonexistent nodes that are not between a pair of nodes will be caught and aborted. Show that the software handles this situation.

15. (A) If a station has several processes, how can a Call Request packet be routed to the correct one?

PART 4
LOCAL AREA NETWORKS

Prior to the development of the microprocessor, most processing was centralized, requiring users to communicate with a single mainframe to which all equipment was attached and upon which their information was stored. The obvious advantage of having a single, centralized computer is that the sharing of resources (such as equipment or information) is a relatively simple procedure controlled by the operating system. For example, a laser printer on a single, centralized computer is accessible to all users through a series of commands that place the object to be printed on the computer's laser printer queue. However, the one great disadvantage of having a single centralized computer is that should it fail, it is impossible for a user to access any of its resources.

The microprocessor has changed the way in which the computer can be used. Now, instead of having all resources available to all users on a single computer, the resources necessary to perform certain tasks may reside on one or more distinct computers. One clear advantage in using a number of computers to perform tasks is that should one fail, users on the others may continue their processing.

However, in a decentralized system consisting of many small, independent computers, it may not be cost effective to supply each computer with the same set of devices (such as laser printers), since any one device may remain idle for long periods. Instead, only a handful of computers may be connected to certain, costly devices. For example, a problem clearly arises when people whose computers are not connected to a laser printer want to print information. One solution is to have the users carry their information on a disk to a computer attached to the printer, or to wheel the printer on a cart from office to office as required. Of course, neither of these approaches

is completely satisfactory, since they can result in a great deal of wasted time on the part of the user.

A similar problem arises when users want to share information. If the information resides on one computer, sharing may be reduced to copying and distributing a disk. This in turn leads to other problems, such as which user will manage the information and how changes to the information will be handled. What is needed is a mechanism whereby individual computers can communicate with others (such as those attached to the laser printer or those with information to be shared).

Although a wide area network (as discussed in Chapter 10) could be used, its cost and speed usually makes it a poor choice for exchanging information among a proliferation of computers within a single organization. This limitation stems from the desire to transfer information between computers at disk (or near disk) speeds to achieve the same speed of service that would be found in a single, centralized system.

To achieve these speeds, a local area network (LAN) is often used to interconnect locally distributed computers. A local area network is similar to a wide area network in that messages are still sent in packets and the stations on the network are identified by an address; however, LANs generally conform to a broad set of requirements, including

Supporting data rates several orders of magnitude greater than those of wide area networks (typically 10 to 100 megabits per second).

Spanning distances of several kilometers (often within a single organization).

Supporting several hundred independent devices.

Offering simplicity of functionality and performance.

Having low error rates, high reliability, and the minimum of dependence upon any centralized components or control.

Permitting efficient use of shared resources (including the network itself).

Maintaining stability under high load.

Allowing fair access to the system by all devices.

Permitting ease of reconfiguration and maintenance.

Being low cost.

Like wide area networks, LANs can be discussed from a variety of different angles, such as topologies, access methods, packet formats, and performance. However, unlike most wide area networks, which exhibit similar characteristics, there is no single generic LAN. For example, the most popular LAN topologies are the bus, ring, and star. Accessing methods are dictated in part by topology, and the performance of different LANs varies greatly depending upon the amount of traffic (the network load).

This part of the book is devoted to the most widespread local area network topologies, notably bus and ring LANs in Chapter 11 and Chapter 12, respectively. In each case, design issues and existing networks are discussed. The Commkit local area network emulators are also examined in detail as a means of showing how the hardware functions.

11
Bus Local Area Networks

11.1
Introduction

A bus local area network is one in which all computing devices (or stations) share a single common communications channel. Stations communicate by passing messages in packets across the network. To distinguish between the different stations, each station on the bus is assigned a unique address. In Figure 11.1, a bus network interconnects five stations.

A typical transmission scenario involves one station transmitting a packet (consisting of at least a message *and* a destination address) one bit at a time onto the network. The bits propagate past all stations and the stations are responsible for reassembling the bits into packets. The packet's destination address is examined by each station: the station with the address that matches the destination address keeps a copy of the packet, and the stations with addresses that do not match the packet's destination address ignore (discard) their copy. A bus network is also known as a broadcast network, since each packet transmitted is received by all stations on the network. Packets are not

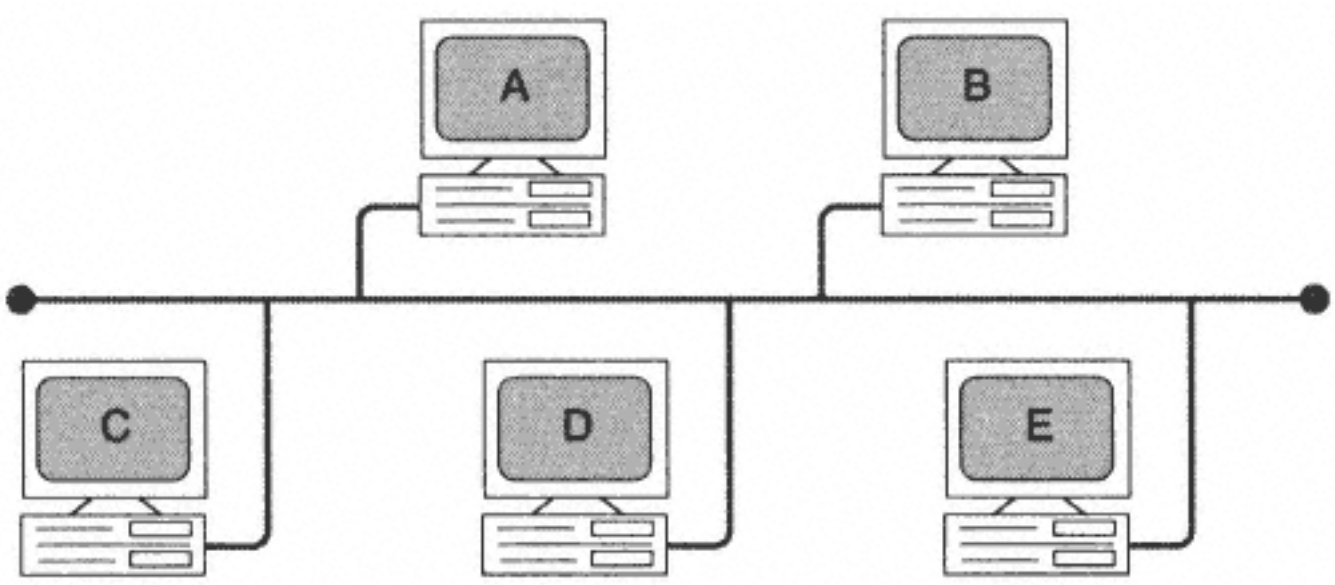

Figure 11.1
A bus network of five stations.

physically removed from the network by any of the stations; instead, each end of the bus has a terminator, which, using resistors, electrically removes the packet.

The bus is a resource shared by all stations on the network and, as such, must support mechanisms that

Prevent one station from monopolizing the network by, for example, transmitting a packet of a size that prevents other stations from transmitting messages.

Resolve conflicts should two (or more) stations attempt to access the network simultaneously.

The first of these mechanisms (preventing a station from monopolizing the network) is usually to restrict the maximum size of the packet and enforce this limit through hardware. However, controlling network access is a somewhat more complex issue and has resulted in a number of solutions.

11.1.1
Controlling Network Access

Controlling access to a bus network is analogous to controlling a group of people at a meeting, all of whom wish to speak. A number of approaches are possible, including

Having one person chair the meeting. The chair is responsible for determining when a person can speak. Once one person has spoken (or signaled the chair that there is nothing to be said), the chair can ask another person to speak.

In a bus network, giving one station the responsibility of controlling network access by polling the others. If a polled station has data to send, the data will be transmitted to the intended destination; otherwise, nothing will be sent. Once the transmission has completed or the controlling device has determined that the polled station has nothing to send (typically through some form of acknowledgment), the controlling device will poll another station. The cycle of polling and transmitting continues indefinitely, permitting each station access to the network.

There are a number of limitations to using a single device to control access to the network. First, should the controlling device fail, none of the other stations can access the network because the polling cycle has stopped. Second, although polling ensures that all stations have equal access to the network, it can result in poor throughput when one station has a large amount of data to send and the other stations have none (since each station must be polled regardless of the amount of data to be sent). Third, a change in the status of the network (such as an existing station leaving the network or a new station joining) requires a mechanism allowing the polling device to determine that a change has occurred.

An alternative to having a person act as the chair of a meeting is to allocate each person a limited time to speak. For example, each person can be permitted to speak for one minutethe first on the hour, the second at one minute past, the third at two minutes past, and so on. The cycle can be repeated once the last speaker has finished. A clock that can be seen by all possible speakers is essential.

A similar approach can be taken to control access to the network, with the exception that each station can transmit only during specific time periods, as dictated by a global clock (as opposed to a centralized device indicating which station can transmit). The underlying assumption is that all stations can synchronize their clocks. For example, at clock tick 1, station 1 transmits, at clock tick 2, station 2 transmits, and so on, until all stations have transmitted, at which point the cycle repeats.

Many of the problems associated with device polling also occur when using synchronized clocks. For example, should a clock on one of the stations fail, the station may start to transmit out of turn; similarly, should one station have a large number of messages to transmit and the others do not, the network will remain idle unnecessarily.

If the money is available, each person can be given a radio that can be tuned to different frequencies. By assigning each person a frequency on which to transmit, all can transmit whenever they feel the need, and those people interested in hearing what someone else has to say can tune into that person's frequency.

Assigning each station a unique frequency band ensures that all stations have simultaneous access to the network. Protocols, such as those used in wide area networks, are required to permit the initial establishment of the call between two stations (using one frequency) and then proceed with the call (possibly on another frequency).

A number of networks have been implemented using this approach, but such networks are costly since they require that each station accessing the network uses hardware that can switch between a number of frequencies. Networks that support multiple frequencies (and multiple simultaneous transmissions) are known as broadband, while networks supporting a single frequency (and thus one transmitting station at a time) are known as baseband.

Although the above networksdevice polling, synchronized clock, and unique frequencyhave been implemented in both commercial and experimental situations, the underlying requirement for some form of centralized control (in the first two networks) and the costs involved (in the third) do not make them as attractive as bus local area networks.

An alternate approach is to resolve the conflict in much the same way humans do when two (or more) people attempt to talk simultaneously: everyone stops and (ideally) one person is allowed to proceed. If a person is talking and simultaneously hears someone else talking, both should stop in

order to determine who should continue. This is resolved in a number of ways (the bigger person, the person with the louder voice, or, more fairly, the person who has remained silent for the longest is allowed to talk first).

A similar algorithm can be applied to stations on the bus network: any station can transmit whenever it has a packet to send and the network is idle (that is, the network is not already in use, indicated by the absence of a carrier signal). Should two stations simultaneously detect that the network is idle, they can both start transmitting, which results in a collision. Collisions can be detected by the station transmitting its packet and simultaneously "listening" to the channel; if the same data is received as is transmitted, there has not been a collision. However, if the data received is not the same as that transmitted, a collision has occurred. The transmit-and-listen approach is used by many bus networks and is given the name Collision Detection (often abbreviated as CD).

Once detected, a collision must be resolved as quickly as possible, since no useful information is being transmitted. Collision resolution, whether it be between people talking simultaneously or between stations on a bus network, is not an instantaneous operation. Thus, there are two main issues to be considered when attempting to minimize the time wasted resolving a collision:

1. How to avoid collisions in the first place.

2. Once a collision occurs, how to determine which station is to transmit next.

Once the network becomes idle, one or more stations may be waiting to transmit their messages. If all of the waiting stations transmit simultaneously (assuming that they all have detected the network being idle), their packets will become garbled, resulting in a collision. To minimize the possibility of collisions when the network becomes idle, and to diminish the length of time the network is idle, there are three classes of algorithm used to determine a station's action when the network is *busy* and the station has a packet to transmit:

Nonpersistent. The station senses the network: if the network is busy, the station waits a random time period before sensing it again. The assumption is that it is unlikely that two stations will wait the same random period, so when the network becomes idle, only one will detect the idle state and start to transmit. However, the nonpersistent algorithm can result in wasted bandwidth if a transmission stops and the network returns to the idle state while the stations with packets to transmit continue waiting for their random time period to expire (see Figure 11.2).

In the example shown in the figure, when the station samples the network and finds it to be busy, it waits for a certain period of time. Meanwhile, the network becomes idle and remains so until the station senses it again. At this point the network is found to be idle and the station can transmit.

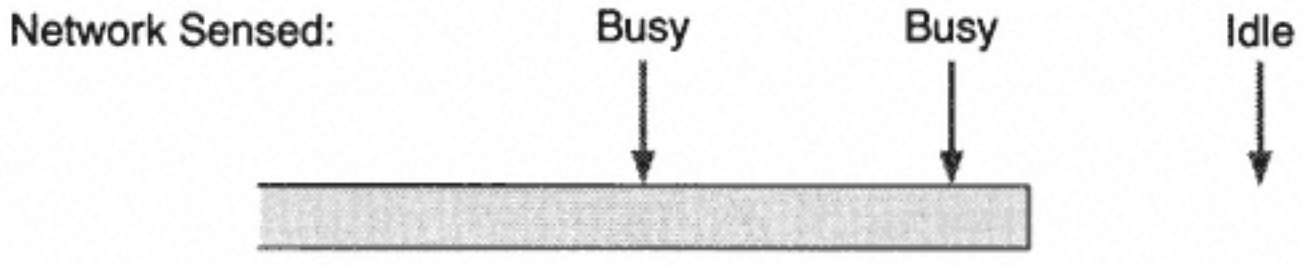

Figure 11.2
Sensing with the nonpersistent algorithm can miss idle periods.

1-persistent. A station with a message to transmit senses that the network is busy and continues to sense it until it is idle, at which point the station starts to transmit its message. The 1-persistent algorithm attempts to ensure that the idle time on the network is as short as possible.

p-persistent. As with the 1-persistent algorithm, the p-persistent algorithm requires that the station continue sensing the network until it becomes idle. At that point, the station will transmit with a probability of *p*; otherwise, it will wait a predetermined length of time before sensing the line again. The *p*-persistent algorithm attempts to minimize the network's idle time while at the same time trying to minimize the chances of stations transmitting simultaneously.

Although all of the above algorithms require that the station sense the bus before transmitting, collisions can still occur if one or more stations detect an idle network and start to transmit. A common form of collision resolution is to have each station back off for a random period of time before attempting to transmit again (somewhat like the nonpersistent algorithm).

Bus networks that require their stations to sense the network before transmitting are commonly known as Carrier Sense Multiple Access (or CSMA). The abbreviation CSMA/CD refers to networks that are carrier sensed and that support multiple access and collision detection. There are a number of features that make CSMA/CD bus networks attractive, including:

The bus is passive. This means that control information such as polling packets is not present on the network; the only network traffic is the data packets sent between stations, and stations are required only to transmit their own packets.

Control on the bus is distributed among all the stations, so the failure of one station does not mean that all communications cease.

11.2 Examples

At present, a number of bus local area networks support CSMA/CD, two of which are discussed in this section: the Ethernet and the IEEE 802.3 bus. A typical bus network controller, designed for both the Ethernet and the 802.3 bus and installed in many PCs, is described at the end of this section.

11.2.1
The Ethernet

Probably the best known CSMA/CD bus network is the Ethernet, a 1-persistent baseband network developed by Xerox in the mid-1970s. The Ethernet is based upon the Aloha system developed at the University of Hawaii-a broadcast network consisting of various stations scattered throughout Hawaii communicating by radio with the university's central computer facilities.

The Ethernet has the following characteristics:

A maximum of 1024 stations can be connected over a distance of up to 2.5 kilometers.

Signals are Manchester (phase) encoded (see Section 6.2.2).

The bus is a coaxial cable.

A data rate of 10 Mbits per second is supported.

Physical Organization

An Ethernet consists of a cable interconnecting a series of stations (devices that contain processors such as PCs, workstations, printers, or disks). The stations connect to the Ethernet by means of a transceiver cable attached to a tap; the tap makes the necessary connections between the transceiver cable and the network. In its simplest form, the Ethernet consists of a single segment; however, more complex layouts consisting of multiple segments are possible, as illustrated in Figure 11.3.

The transceiver cable allows stations to be placed up to 50 meters away from the Ethernet. To avoid interference between station taps, stations must be a minimum of 2.5 meters apart (hence the maximum distance of 2.5 kilometers for 1024 stations). The maximum recommended segment length is 500 meters. The 50-ohm coaxial cable has a diameter of 0.4 inches.

The Ethernet topology is considered a loopless "unrooted tree," which

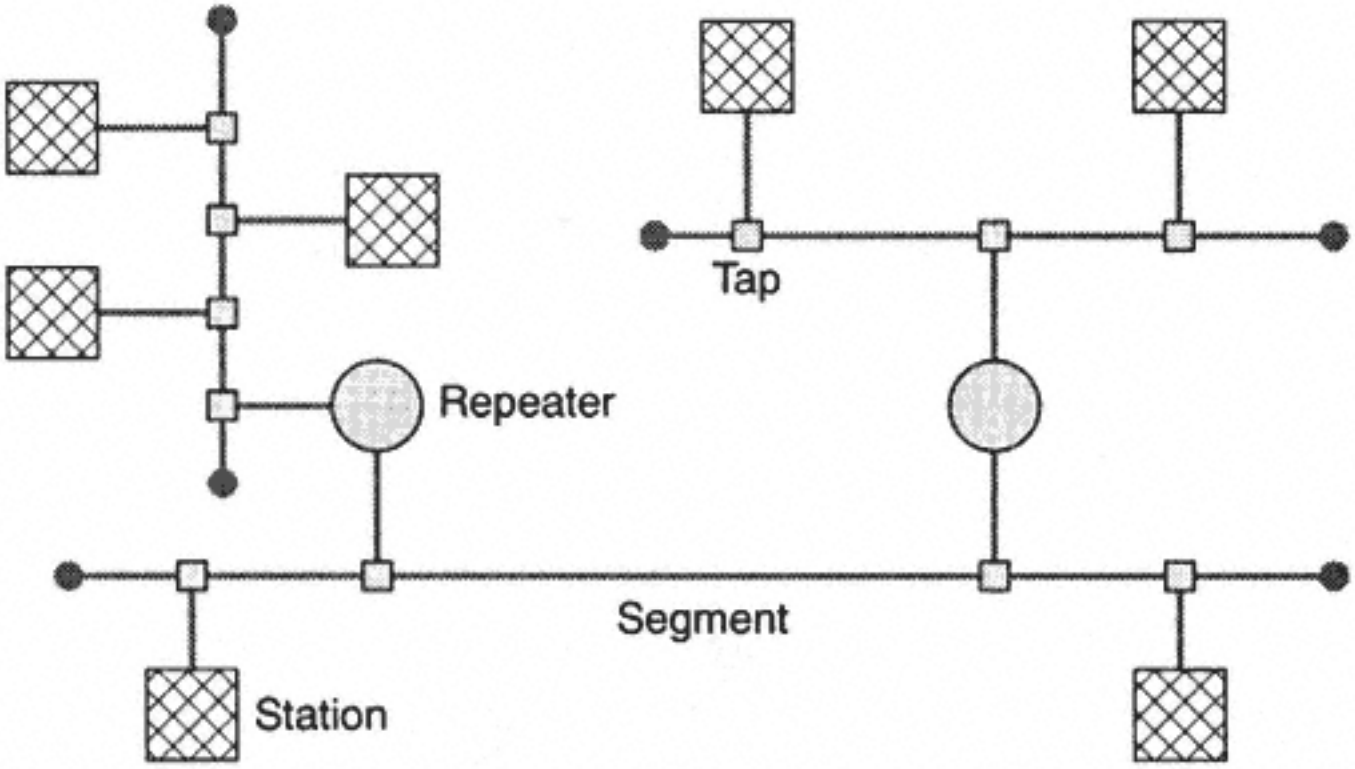

Figure 11.3
An Ethernet with multiple segments.

means that a packet traverses the entire tree and individual branches cannot be used independently.

The Ethernet Packet

All data is sent across the Ethernet in an Ethernet packet. The start-of-packet is indicated by a synchronizing pattern of 64 bits of alternating 1s and 0s, ending with two consecutive 1s. All bits are Manchester encoded. The packet structure, shown in Figure 11.4, consists of five fields (the numbers in parenthesis indicate the size of a field, expressed in bytes).

The fields in an Ethernet packet are defined in the following paragraphs.

The fields DST and SRC are, respectively, the destination (i.e., the station for which the packet is intended) and source (i.e., the station sending the packet) addresses of the packet. All stations on the network are assigned a globally unique 48-bit station address (supplied by Xerox, thereby ensuring each address's uniqueness). This address is inserted into the Ethernet packet as the source address (SRC).

The destination address (DST) is also 48 bits long and can be one of the following:

A unicast address, identifying a single destination on the network and unique to the station. A packet sent with a unicast address is discarded by all stations except the one with the address matching that in the DST field. The source address (SRC) is a unicast address.

A broadcast address, identifying all stations on the network and shared by them. A packet sent with a broadcast address is kept by all stations.

A multicast address, identifying a group of stations on the network. The number of stations sharing a multicast address is typically determined by the application and can vary over time. A packet sent with a multicast address is kept by those stations sharing the address.

Table 11.1 shows how addresses are distinguished by their high-order bit.

All commercially available Ethernet hardware supports at least the station's unicast address and the broadcast address. Multicast addresses are also supported by some Ethernet hardware to a limited degree. For example, Digital's DEUNA (Digital Equipment UNIBUS Network Adaptor) hardware can identify up to 10 distinct multicast addresses, while the Intel 82586 can filter up to 64. However, the multicast address recognition algorithm used in the 82586 can result in multicast addresses being accepted by the hardware that are not supported by the station.

Finally, some Ethernet hardware permits stations to enter promiscuous

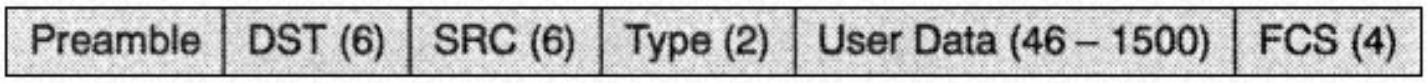

Figure 11.4
The Ethernet packet structure.

Table 11.1 Ethernet address structures.

Identifier	High-Order Bit	Remaining 47 Bits
Unicast	Zero	Zeros and ones
Broadcast	One	All ones
Multicast	One	Zeros and ones

mode, in which all packets on the network, regardless of the value of the destination address, are made available to the station.

By itself, the Ethernet offers only a point-to-point datagram service between *stations*; the destination address offers no indication of the intended destination process, application, or protocol. The two-byte Type field overcomes this limitation by allowing the Ethernet software to route the User Data to a specific higher-layer protocol, application, or process. For example, a station may support two services, each associated with its own Type value. An incoming packet is routed to the service indicated by the value of the Type field.

The User Data field contains the data sent from one station to the station(s) specified by the destination address. It cannot contain less than 46 or more than 1500 bytes. All data is transmitted transparently; that is, all possible byte values are considered valid.

The minimum packet size has been chosen to ensure that all collisions will be detected, regardless of the length of the Ethernet (up to its physical maximum of 2.5 kilometers). For example, in Figure 11.5, stations A and B each transmit a short packet simultaneously; the collision is undetected by the transmitting stations because the packets are too small. Similarly, a station will miss the collision if its transmission stops before the station receives a bit from any other packet. This problem is illustrated in Figure 11.6.

From the figure, one can conclude that a collision will be detected only if a station transmits for the length of time it takes a bit to traverse *twice* the length of the Ethernet. The Ethernet transmits 107 bits (10 megabits) a second, or one bit every 107 of a second. Since the information travels at the speed of light (108 meters a second), the "length" of one bit is 10 meters. Therefore, 500 bits fill twice the maximum length of an Ethernet

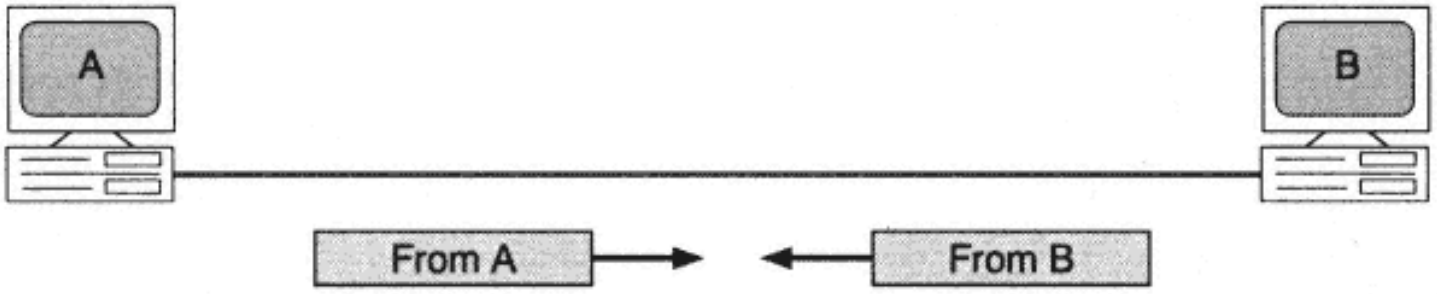

Figure 11.5
A collision cannot be detected, since the packets are too small.

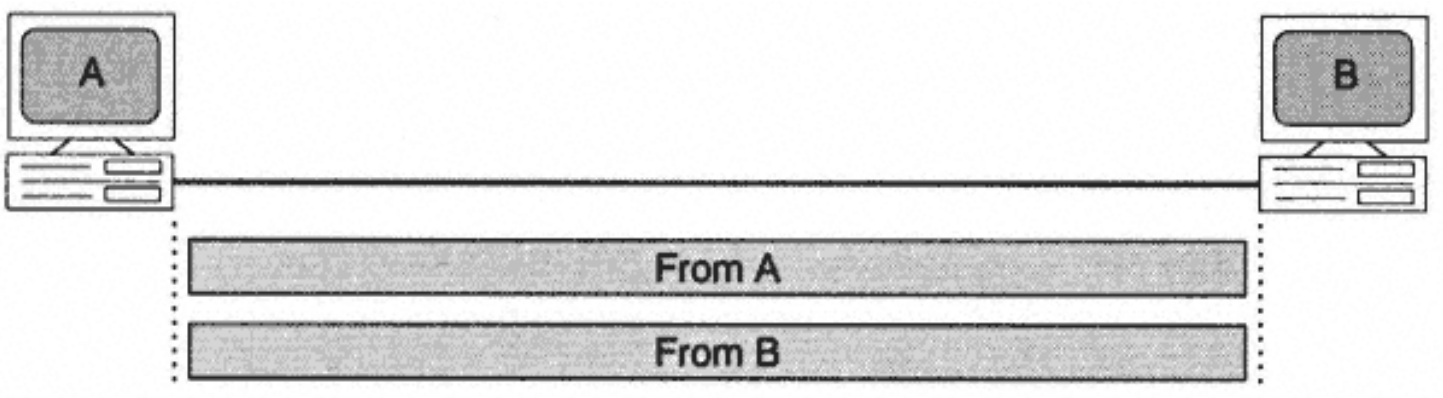

Figure 11.6
A collision cannot be detected, since the packets are still too small.

(5 kilometers, or 5000 meters). By rounding the number of bits to 512 (a power of 2), the minimum number of bytes in a packet becomes $\frac{512}{8}$, or 64. Four fields within the packet have fixed sizes (DST, SRC, Type, and FCS), amounting to 18 bytes. This means that the minimum size of the User Data field is 64 18, or 46, bytes!

The maximum User Data field length ensures that a transmitting station does not monopolize the Ethernet for an undue period of time. Most Ethernet hardware supports anti-jabber features, which cut off the transmitter if the packet's length exceeds the maximum number of bytes permitted (the total number of bytes from all fields).

The Frame Check Sequence (FCS) field is a 32-bit CRC value covering all bits in the packet (excluding the preamble). The FCS transmitted is included in the FCS received by each station once the Ethernet has gone idle (see below). If the calculated FCS is equal to a predefined value, the packet is assumed to be correct.

The end-of-packet is determined when no traffic is detected on the Ethernet. To ensure that each station's receiving hardware has time to examine the packet and supply it to a higher level of software, there is a minimum packet spacing of 9.6 microseconds.

When collisions occur, the stations involved are expected to wait random time periods before attempting to send their packets again. The use of a random delay is based upon the assumption that if the stations don't pick the same random value, the collision will be avoided. The Ethernet implements a truncated binary exponential backoff algorithm that each station uses if a collision occurs. The algorithm works as follows: each time a collision occurs, the station must wait a time period based upon a random number (see below) multiplied by the length of time it takes for a bit to travel from one end of the network to the other (51.2 μseconds) before attempting to transmit again. The random number is chosen as an integer in the range $0 \leq r \leq 2k$, where r is the random number and k is the number of attempted transmissions or 10whichever is less; the growth is truncated at $k = 10$ (hence the name of the algorithm).

The station aborts its transmission after 16 collisions.

11.2.2
IEEE 802.3

In the early 1980s, the IEEE set up a local area network standards committee known as IEEE 802. Part of the mandate of the 802 committee has been to define Data Link and Physical Layer specifications for LANs. All networks are uniquely defined in terms of a Physical Layer and a MAC (Media Access Control) Layer. The MAC Layer overlays the Physical Layer and part of the Data Link Layer.

One of the networks proposed by the 802 committee is 802.3, a CSMA/CD bus network based upon the Ethernet. This standard differs from the Ethernet standard in three main areas: electrical connections, addressing, and the Type field.

802.3 allows both 16-and 48-bit addresses (the 10-Mbps standard requires 48-bit addresses). Unicast, multicast, and broadcast addresses are all supported. IEEE also assigns each station its own global address, which allows the station to be uniquely identified anywhere in the world.

The 802.3 packet structure is essentially the same as the Ethernet packet structure, with the exception of the Type field (all sizes are expressed in bytes), as shown in Figure 11.7.

In 802.3, the Type field is replaced by a two-byte Length field, which indicates the length of the User Data field, expressed in bytes. The User Data field can be from 46 to 1500 bytes long, although the Length field can have values less than 46. (If the message is less than 46 bytes long, it must be padded to the minimum length.)

Note that it is possible for both Ethernet and 802.3 packets to co-exist on either network. For example, Ethernet packets can be transmitted on an 802.3 network as long as the value in the Type field does not fall in the range of possible Length values. This is achieved by setting the value of the Type field to anything greater than the maximum Length value (i.e., greater than 1500).

The 802.3 protocol has a number of variants that are expressed in the form *xyz*. where *x* is the line speed expressed in megabits per second, *y* indicates the type of physical medium, and *z* is the maximum length of a segment, expressed in hundreds of meters. The original 10-megabit baseband 802.3 implementation on thickwire coaxial cable (500-meter segments) is written as 10Base5. Other variants include 10Base2 (thinwire coaxial cable, more commonly known as Cheapernet), 1Base5 (1-megabit over twisted-pair), and 10BaseT (10-megabit over twisted-pair).

The 802.3 CSMA/CD network also implements the truncated binary exponential backoff algorithm to handle collisions.

Other 802 protocols will be discussed further in subsequent chapters.

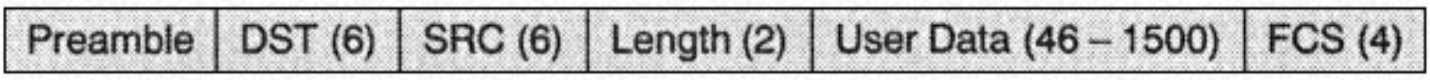

Figure 11.7
The 802.3 packet structure.

11.2.3
Network Access Hardware

Commercially available network access hardware for the Ethernet and 802.3 typically consists of three distinct chips; the relationship of the chips is shown in Figure 11.8.

The transceiver chip is responsible for transmitting and receiving the bits. A transceiver such as the Intel 82502 performs additional functions, including

Monitoring the number of bits the station has transmitted and cutting off access to the Ethernet if the number exceeds the maximum allowed in a packet (this is the anti-jabber function). The station is barred from transmitting for about 420 milliseconds, after which the transceiver again permits access to the network.

Comparing the bits read off the network with those being transmitted. Should the bit sent not equal the bit received, a collision has been detected. This information is relayed back to the serial interface.

All signals received by the transceiver are Manchester encoded; they are not decoded but rather passed directly to the serial interface. The serial interface, such as the Intel 82501, extracts the clocking signal from the incoming data stream and generates the equivalent internal bit value. It is also responsible for encoding the outgoing data with a 10-MHz Manchester signal supplied to the transceiver. If a collision is detected (by the transceiver), the condition is forwarded to the LAN coprocessor.

The LAN coprocessor is a separate processor that implements the Ethernet protocol and performs the following functions:

Calculates the Frame Check Sequence (FCS) for outgoing and incoming frames.

Supports unicast, broadcast, and up to 64 multicast addresses.

Transmits Ethernet packets supplied from the CPU.

Filters and supplies Ethernet packets to the CPU.

In the Ethernet (802.3 10Base5, 10 Mbps, baseband, 500m segments), the LAN coprocessor and serial interface can be on the DTE itself, connected to the transceiver by a transceiver cable. The transceiver physically taps the LAN. In the Cheapernet (802.3 10Base2, 10 Mbps, baseband, 200m segments; and 10BaseT, 10 Mbps, baseband, twisted-pair), the three chips are typically on a single board, with the transceiver attached to a coaxial

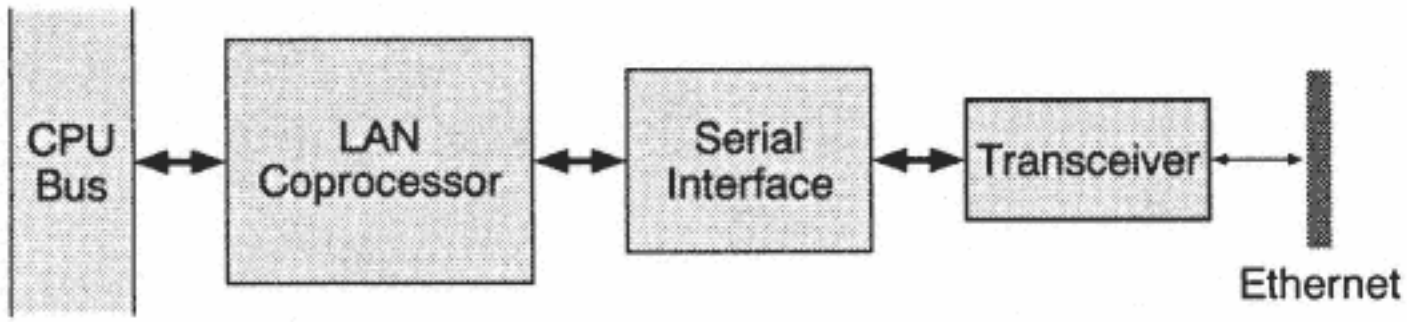

Figure 11.8
Typical 802.3/Ethernet hardware.

cable connector. The LAN connects to the transceiver via a bayonet or BNC T-connector ("BNC" is an acronym of the device and its inventor: Bayonet, Neil Concelman).

The Intel chipset described above supports the entire 802.3 CSMA/CD MAC functions. Many other Ethernet/IEEE 802.3 chipsets are available from manufacturers, including Advanced Micro Devices (AMD), National Semiconductor, and Seeq Technology.

11.3
The Commkit Bus Network

The communication hardware available with a standard PC does not support any type of bus network. Ethernet cards (using controller chips such as the Intel 82586) are commercially available and can cost upwards of several hundred dollars. However, it is possible to write software that emulates a CSMA/CD bus network by using the PC's serial ports. Commkit is supplied with the following bus network software:

`buslow.cthe` C routines that allow the PC to emulate a bus network.

`busdefs.hthe` header file containing the various data structures used by `buslow.c`.

`busex.ca` set of processes that allow messages to be entered onto the network from the keyboard.

`busex.hthe` bus network packet structure and return codes used by `busex.c`.

11.3.1
Hardware Requirements

The Commkit bus network software requires a minimum of two PCs, each with two serial ports (assumed to be `SP1IH` and `SP2IH`) and sufficient null-modem cables to allow the PCs to be linked together in a bus-like fashion. Figure 11.9 shows a Commkit bus network of four stations. The network consists of four PCs, each supporting two serial ports, with an RS-232 null-modem cable connecting pairs of ports (except the outermost two, which must *not* be connected). There is no prescribed ordering of serial ports (that is, any port can connect to any otherfor example, 1 to 1, 1 to 2, or 2 to 2).

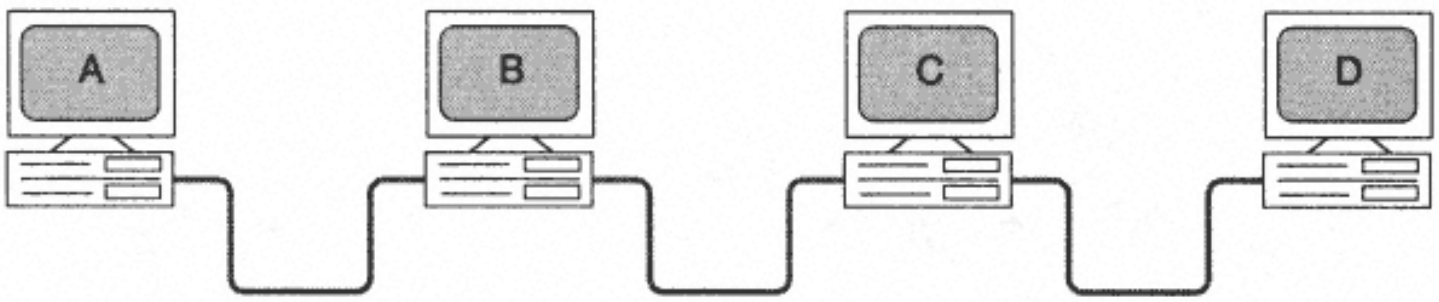

Figure 11.9
The configuration of a 4-station Commkit bus network.

11.3.2
Design

The bus network supplied with Commkit is truly CSMA/CD in that before a packet is transmitted the network is first "sensed"; if it is idle, the packet is transmitted. Bytes within a packet are sent one at a time out of each of the PC's serial ports. In Figure 11.10, station B transmits a packet *N* bytes in length across a 4-station Commkit bus network.

When a nontransmitting station receives a byte on a serial port, a copy of the byte is taken and the byte is then forwarded out the "other" serial port. The bytes are assembled into packets for subsequent processing.

A collision is detected by the transmitting station receiving a byte. Similarly, a receiving station can detect a collision if it receives bytes on both ports. As in a true CSMA/CD bus network, transmissions are aborted if a collision is detected.

The Packet Structure

Each message that is transmitted on the network is sent in a packet. In addition to the message, the packet contains the address of the intended destination station, the address of the transmitting station, and two CRC bytes.

Packets are transmitted transparently; that is, their message portion can contain any byte value (i.e., from 0 through 255). When transmitted, the packet is delimited by the packet-framing characters *STX* and *ETX*. Any bytes within the message having values *ETX* or *DLE* are prefixed by *DLE* when the packet is transmitted. The *DLE* prefix bytes are removed as the packet is received. They are not included in the total message size nor are they made available to the destination process. Neither the destination address, DST, nor the source address, SRC, is prefixed by *DLE*; since the position

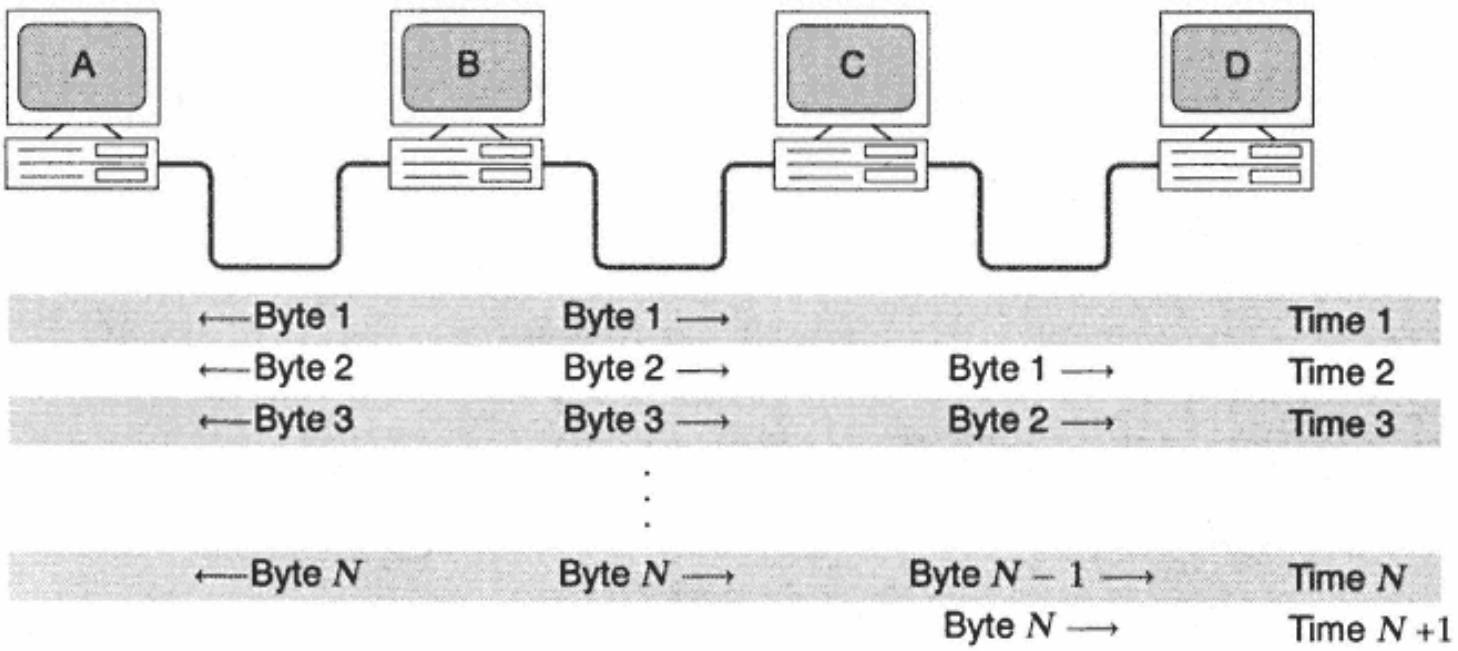

Figure 11.10
Transmitting a message on the Commkit bus network.

of these bytes is well-known, there is no chance of confusing them with the trailing *ETX*.

Each byte in the packet (excluding the packet-framing bytes, *STX* and *ETX*. and any prefix *DLE* bytes) is included in the 16-bit CRC, which is sent as part of the packet. Should either of the CRC bytes be an *ETX* or a *DLE*. it is also prefixed by a *DLE*. The CRC bytes are removed from the packet and are not made available to the destination process. A packet received is assumed to be error-free if the receiving station's calculated CRC has a value of 0 when the end-of-packet *ETX* is found. The overall packet structure is shown in Figure 11.11 (the number of bytes in a specific field is indicated in parenthesis).

The internal packet structure is defined in `busex.h` as

```
struct packet
{
char net_dst;
char net_src;
char message[PKT_MAX + 2];     /* Includes two
                                  CRC bytes */
};
```

The maximum size of the message is defined by `PKT_MAX` in `busex.h`. `PKT_MAX` can be changed, but should not be allowed to exceed the size of a Commkit message (128 bytes) or be less than the minimum packet size.

To ensure that all stations have received at least one byte while the transmitting station is still transmitting, there is a minimum packet size defined by `MIN_LEN` in `buslow.h`. The choice of minimum packet length is determined by the number of cables connecting the stations on the network. For example, if the network consists of *N* cables, the minimum packet length is $2 \times N$. With $2 \times N$ as the minimum packet size, collisions can be detected by all stations on the network.

The Bus Network States

At any moment, the bus network can be in one of three states: idle, transmitting, or receiving.

If the station has not detected any traffic on the network, and there are no packets to be transmitted, it is in the idle state. When the station is first initialized, it is idle. It is also idle after a packet has been transmitted or received, or a collision has been detected.

Figure 11.11
The Commkit bus frame structure.

Either of the following conditions will cause the station to leave the idle state:

1. A packet is made available for transmission (the station's state is changed to transmitting).

2. A byte is received on either of the serial ports (the station's state is changed to receiving).

The transmitting state is entered only after the time delay associated with network access has expired. There are two sources of time delay: one, each time the transmitting station sends a packet that collides with another packet, access is denied for progressively longer and longer periods; two, once a packet has been received or transmitted, access to the network is delayed by `PKT_DELAY` clock ticks in order to allow other stations access. As soon as the time delay reaches 0, the station can resume sensing the network.

The transmitting state is entered from the idle state (i.e., no traffic has been detected on the network) if either of the following conditions arise:

1. A pending packet (i.e., a packet whose transmission was aborted because of a collision) is available for transmission.

2. A new packet is available for transmission.

To ensure that the packets are transmitted in the same order in which they are supplied to the bus network emulator, any packet that has been delayed by a collision has priority over new messages. This priority is achieved by having the message queue associated with the bus network read *after* the transmission of a packet.

Message transmission begins in the idle state with the transmission of two *STX* bytes, one out each serial port. The contents of the packet data structure are then sent, one byte at a time, out each serial port. Each byte from the packet is included in the CRC prior to being transmitted; *DLE* insertion also takes place at this point. Once the message bytes have been transmitted, the two CRC bytes are sent (with *DLE* insertion if necessary) followed by an *ETX*. The transmitting station then returns to the idle state. The state diagram for a transmitting station is shown in Figure 11.12 ("`Send()`" refers to putting bytes on the network).

Ideally, only one station at a time will be transmitting on the network. However, on occasion, two (or more) stations may be in the idle state when a packet is ready to be transmitted, resulting in a number of stations commencing a transmission at the same time. The stations will continue to transmit until they receive a byte, at which point a collision is detected (a station can both receive and transmit through the same port because the channels are physically separate). Figure 11.13, for example, shows the effect on the sample bus network if stations A and D are to transmit simultaneously (a "*" denotes that the collision has been detected and "⇔" indicates that two bytes are on the cable simultaneously).

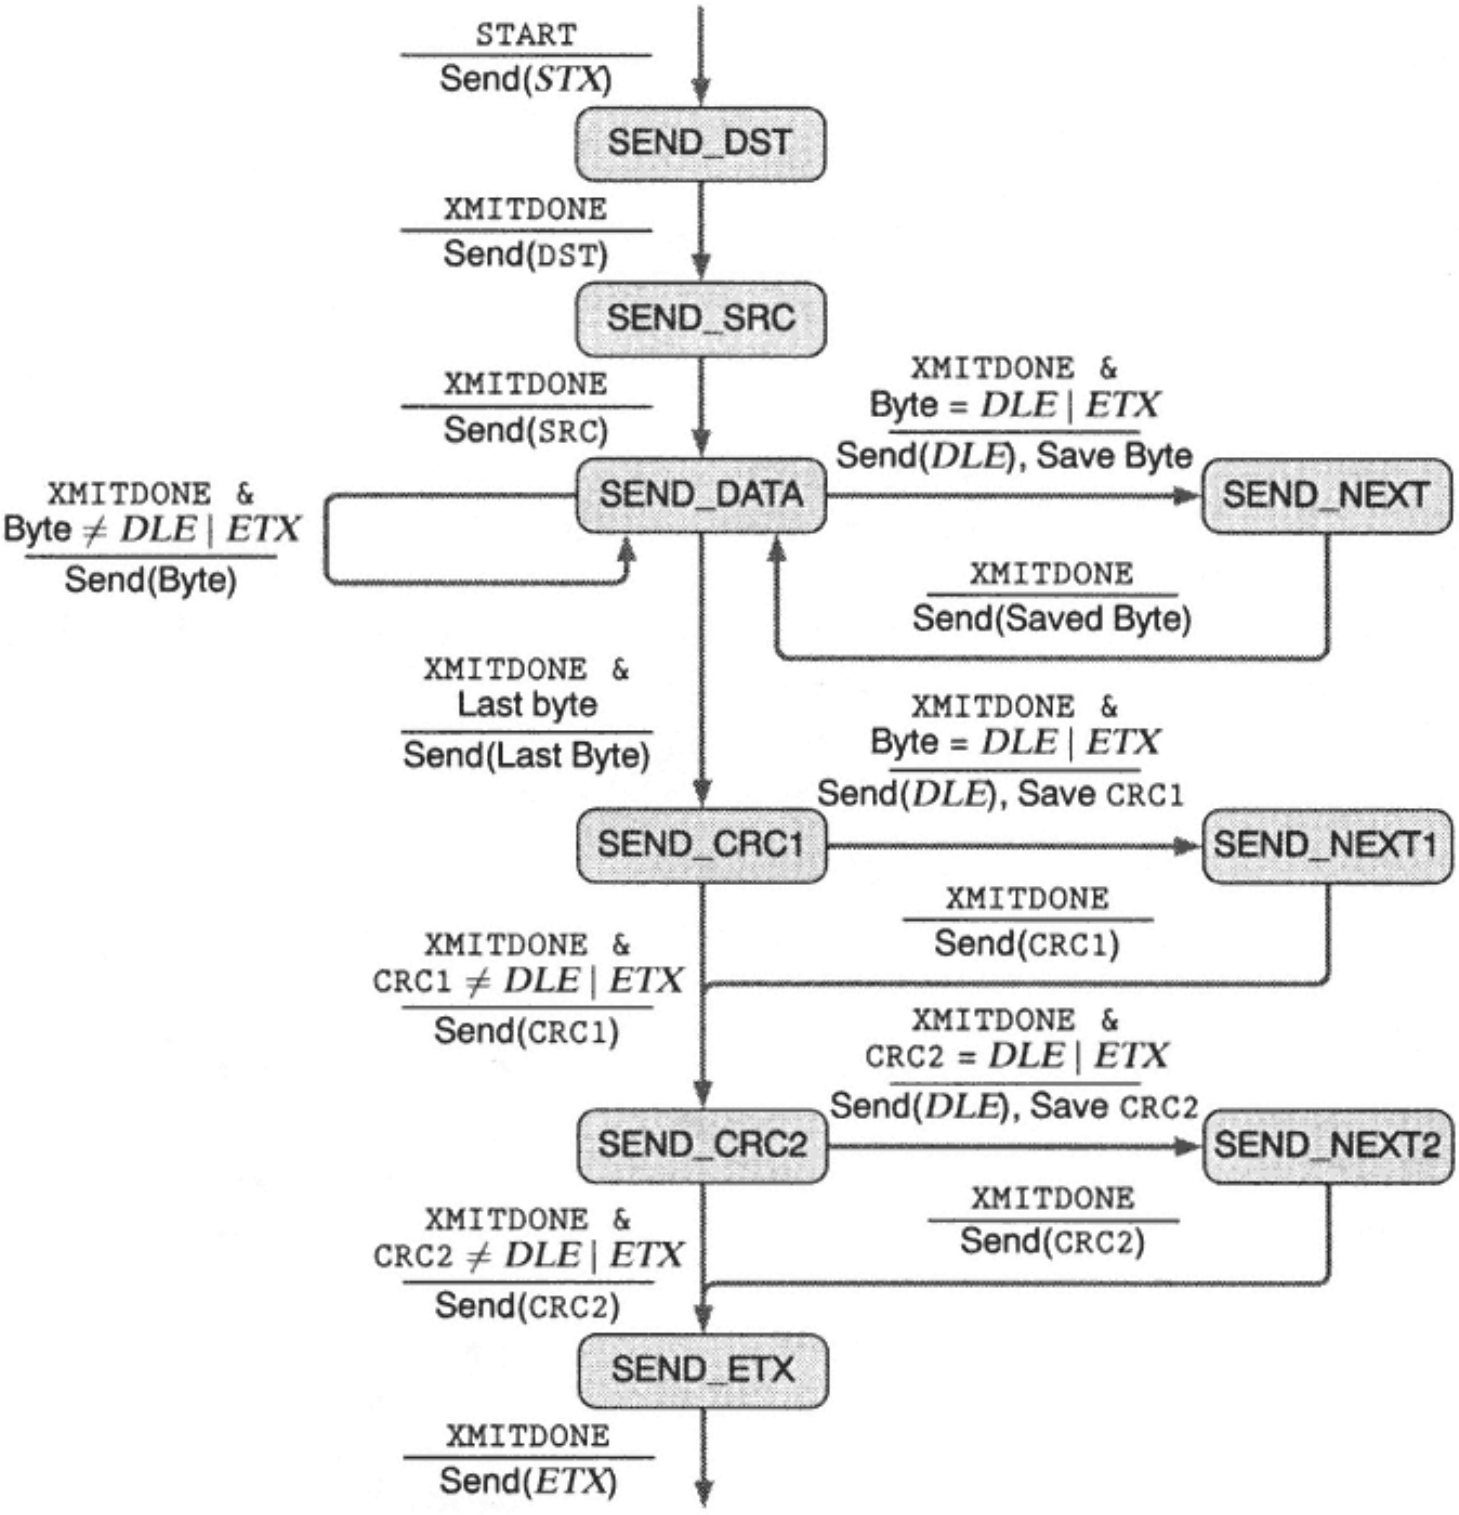

Figure 11.12
Bus network transmission state machine.

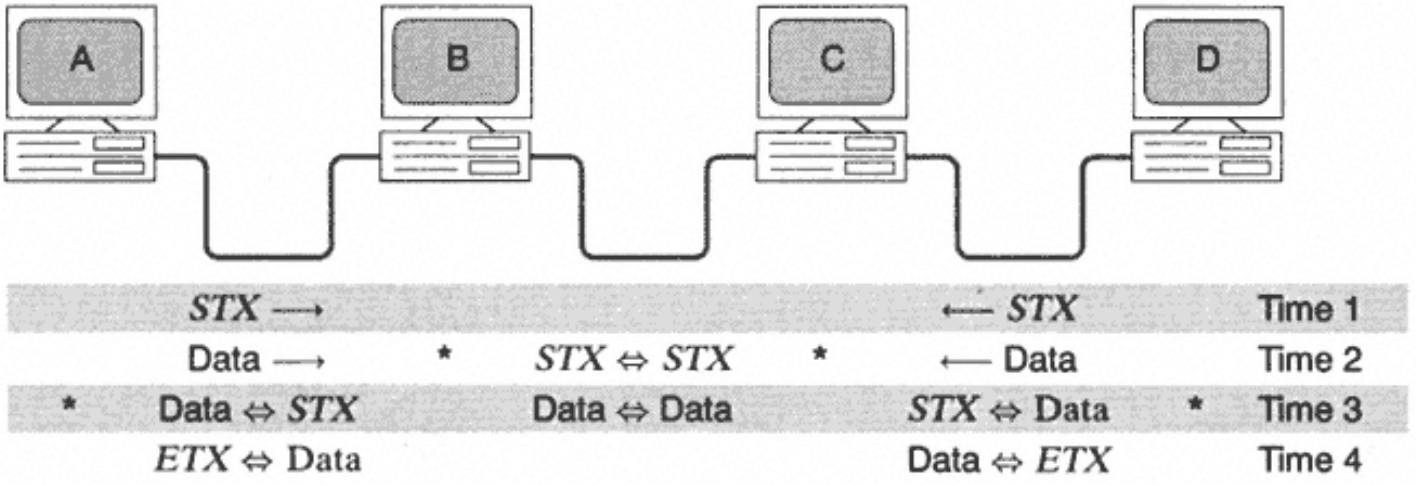

Figure 11.13
Simultaneous transmission and subsequent collision detection.

Upon detection of an incoming byte, the transmitting station immediately aborts its transmission and sends out an *ETX* on both ports. If a *DLE* was the last byte sent, another byte is sent prior to the transmission of the *ETX*. thereby avoiding the possibility of generating a *DLEETX* sequence rather than simply an *ETX*. Each transmitting station then backs off for a period of time determined by its station identifier and the number of collisions associated with the message.

The station then reenters the idle state and cannot transmit until the backoff period has expired. When the station detects that the delay has ended, the pending message can be sent as soon as the network becomes idle.

The receive state is entered when an *STX* has been detected on either of the serial ports and the station is in the idle state. The receive software forwards the *STX* out the serial port that did not receive the byte; all subsequent bytes associated with this packet are expected to be received on the same serial port (port 1 or port 2, *not* both). See Figure 11.14 for an example.

All bytes following the *STX* are copied into the packet data structure as well as being forwarded through the outgoing port. The first two bytes following the *STX* are treated as the destination and source address bytes, respectively. Subsequent bytes are stored in the next available location in the packet data structure. *DLE* prefix bytes are also forwarded, but are not kept in the message buffer. All bytes stored in the packet are included in the CRC calculation. Upon reception of the end-of-packet indicator (*ETX*), the CRC is checked. If its value is 0 and the destination address byte matches that of the station, the packet is sent to the process responsible for handling incoming messages. In all other cases, the message is discarded and the bus network returns to the idle state. Figure 11.15 illustrates the state machine for a station receiving a packet.

Normally, all incoming bytes are received on the same port. However, should a byte be received on the outgoing serial port, a collision has occurred. Upon detection of a collision, the receiving station no longer stores the incoming bytes in the packet buffer; instead, they are discarded. To ensure that all stations on the network can detect the collision, all bytes (received on either port) are forwarded out of the serial port opposite to the one on which they were received (receive on 1, forward on 2; receive on 2, forward on 1), as illustrated in Figure 11.16. The receiving station continues to forward the bytes until an *ETX* is detected on both ports, at which point the station enters the idle state.

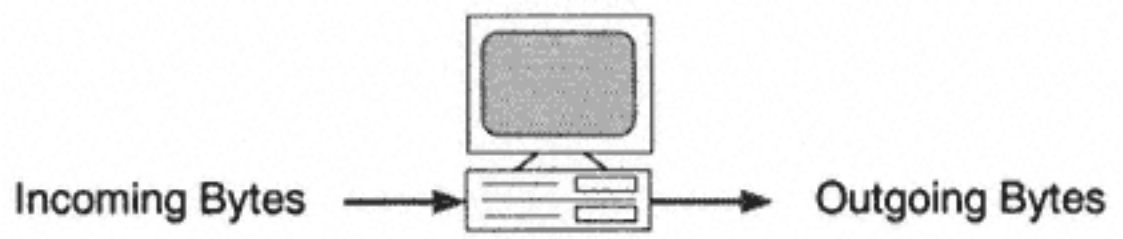

Figure 11.14
Incoming bytes are expected on a single port.

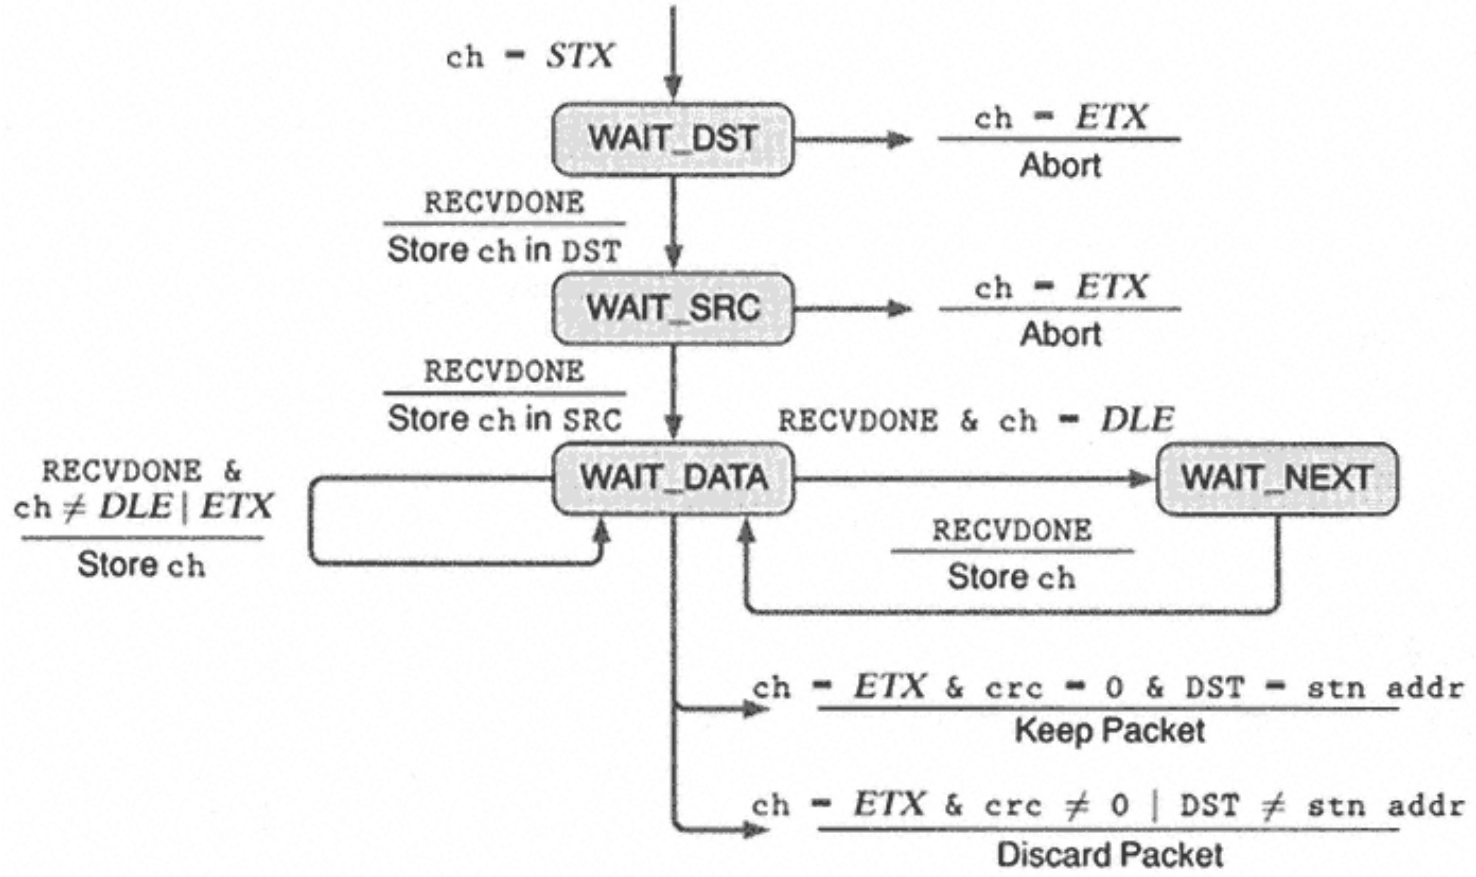

Figure 11.15
Bus network receive state machine.

Note that the collision-handling technique described for reception will not necessarily be recorded as a collision on all receiving stations. For example, in the network of four stations depicted in Figure 11.17, if stations A and D start to transmit simultaneously, the receiving stations B and C will both detect a collision, since they receive data on both ports (a "*" denotes that the collision has been detected).

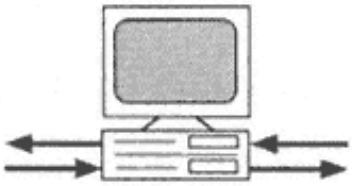

Figure 11.16
Bytes received on one port are forwarded on the other port.

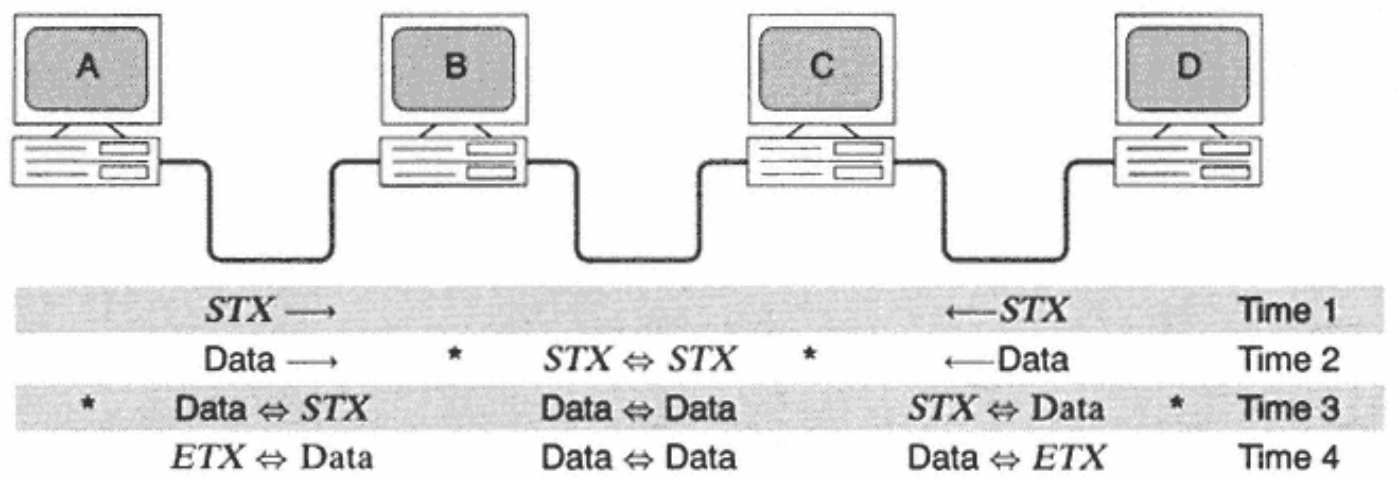

Figure 11.17
Detecting collisions (part 1).

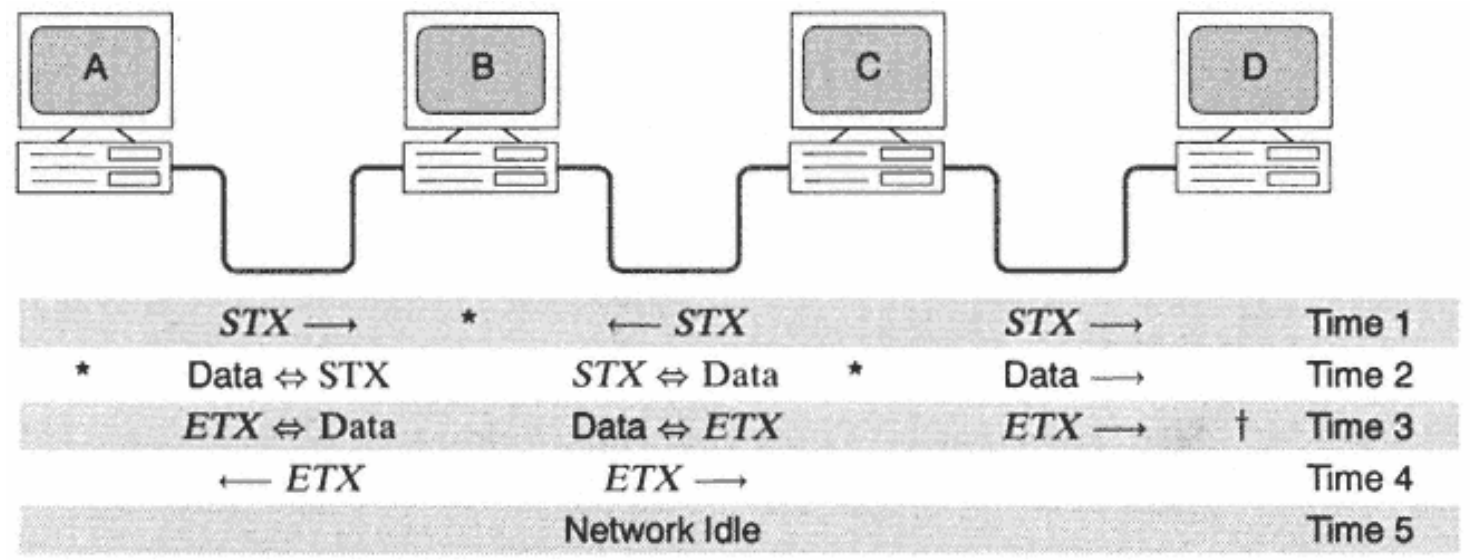

Figure 11.18
Detecting collisions (part 2).

However, should stations A and C transmit simultaneously, only station B (of the two receiving stations B and D) will detect the collision as a true one, since it receives data on both serial ports. Station D, on the other hand, receives only bytes from station C (since C will not forward the bytes from station A and station D cannot distinguish between the two stations' bytes, anyway), with the result that during a collision station D receives a packet with a CRC error from station C. In both situations, the packet is discarded but is recorded as a collision only by stations between the two transmitters and as a CRC error by all other receiving stations; this is illustrated in Figure 11.18 (a "*" denotes that the collision has been detected; a "+" indicates that the collision has been detected as a CRC error).

11.3.3
The Bus Network Implementation

The Commkit bus network software has been written in a layered fashion to give different high-level applications access to the bus emulation software. The overall structure for the example presented in this section is given in Figure 11.19.

The processes found in the 1-line message-handling software (`busex.c`) communicate with the bus emulation software (`buslow.c`) using `send()s` and `recv()s`. The bus network emulation software `buslow.c` requires `commkit.obj`, `srcalls.obj`, and `crc.obj` in order to function. `Commkit.obj` calls the bus network emulation software via the entry point

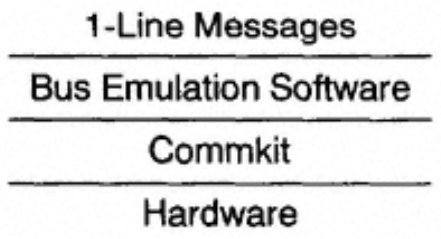

Figure 11.19
The layering of the bus
network example.

Table 11.2 The bus network implementation of `low_level()`.

source	action	ch
CLKIH	Not supplied	Not supplied
SP1IH	RECVDONE	Character Received
SP2IH		
SP1IH	XMITDONE	Not supplied
SP2IH		

`low_level()`. `Low_level()`, in turn, calls the bus network software via a common entry point, `bus_emulator()`; the cycle is shown in Table 11.2.

Network Emulation Software

The three states described in Section 11.3.2 are implemented directly in `bus_emulator()`.

If the station is neither transmitting (sending a message) nor receiving (accepting and forwarding a message), it is in the idle state (indicated by `bus.state` having a value of `IDLE`). While in the idle state, the bus emulator can be called by one of three interrupt handlers: the clock, `CLKIH`; serial port 1, `SP1IH`; or serial port 2, `SP2IH`.

Interrupts from `CLKIH` are used for a number of purposes. First, if `bus.delay` (the delay required before this station is allowed to transmit) is non-zero, it is decremented on each tick; when 0 is reached, the station is allowed to enter the transmission state, `TRANSMIT`. This state is entered if there is a message pending (indicated by `bus.pending`) or if there is a message on the `BUSNET` queue waiting to be sent (the `BUSNET` queue is actually `SP1IH's` queue). Procedure `send_start()` is called when a message is to be sent; `send_start()` transmits an *STX*. which is sent out both serial ports. `RECVDONE` interrupts from either `SP1IH` or `SP2IH` cause the emulator's state to change to `RECEIVE` if the incoming character is an *STX* (see `recv_start()`). The emulator "remembers" the incoming and outgoing ports (for collision detection) by storing the specific port identifiers in `bus.inport` and `bus.outport`, respectively. Forwarding to the next station is by supplying the character to the port on which it was not received.

The receive state is entered when `bus.state` has a value of `RECEIVE`. `Recv_byte()` is called whenever a character is received on either serial port (`SP1IH` or `SP2IH`). Under normal circumstances, the `recv_byte()` state machine determines what to do next with the incoming character (see Section 11.3.2). Once an entire packet has been received and the CRC determined to be valid, the message is put on the queue `BACKGROUND_1`. However, if the port on which the incoming character has been received is not

the same as the original reception port, a collision has occurred, causing the reception to be aborted (signaled by `bus.collision` set to `TRUE`). After a collision is detected, reception continues until an *ETX* is received on both ports.

Clock interrupts still occur in the `RECEIVE` state and are used to ensure that the loss of communications (for example, a missing *ETX* because of the transmitting station crashing) does not leave the receiving station waiting "forever" for the missing byte. A watchdog timer, `bus.xmit_watchdog`, is incremented on each clock tick and reset whenever a byte is received. Should the flow of bytes cease, the timer value increases until it exceeds a maximum, causing `recv_abort()` to be called. `Recv_abort()` puts the station back to the idle state and sends a `RECV_ABORTED` message to `APPLICATION`.

Each station also supports promiscuous mode, in which any message received with a valid CRC is accepted regardless of its destination. Promiscuous mode can be specified by assigning `TRUE` to the global variable `promiscuous` after Commkit has been initialized. By default, `promiscuous` is `FALSE`.

The transmit state is entered when `bus.state` has a value of `TRANSMIT`. An `action` value of `XMITDONE` results in the next byte in the packet being sent. Since there is no guarantee that both ports will generate interrupts at the same rate, transmissions are controlled separately in procedure `send_byte()`. When both ports have transmitted the final *ETX*. the message `MESSAGE_SENT` is forwarded to the transmitting process and control returns to the idle state.

However, an `action` value of `RECVDONE` indicates that another station is attempting to transmit at the same time as the first station, meaning that a collision is occurring. When a collision has been detected, `send_collision()` is called and `bus.collision` is set to `TRUE`. Thereafter, the `send_byte()` software sends an *ETX* (or a byte followed by *ETX* if a *DLE* was previously sent). `Bus.state` returns to `IDLE` when an incoming *ETX* is found.

`Collision_action()` determines how long the station must wait before attempting to transmit again. A simplified version of binary exponential backoff is used: the station identifier, `bus.stn_id`, is used as the random number between 1 and 4 (stored in `bus.backoff`), which is then multiplied by the number of times the station has collided while trying to transmit (indicated by `bus.backcount`) and the packet delay, `PKT_DELAY`. If the number of collisions exceeds the upper limit (defined in `BKOFF_LIMIT`), the station gives up trying to transmit this message and informs the transmitting process.

The clock interrupt is used as a watchdog to determine whether the other station involved in the collision has sent an *ETX*. When a collision has been detected, `bus.delay` is set to a non-zero value; if an *ETX* is not found by the time `bus.delay` reaches 0, the emulator returns to the idle state and a `MISSING_ETX` message is sent to the transmitting process.

One-Line Message Software

By itself, the bus network emulation software described in the previous section handles only the transmission and reception of packets: packets sent are taken off the BUSNET queue and packets received are put on the BACKGROUND_1 queue. To permit experimentation with the bus network, Commkit includes software that allows the user to enter a packet (destination and source address and data) from the keyboard. Packets are sent by bus emulation software across the network to the station specified in the destination address. Packets received by emulation software for the user's station are displayed locally.

The 1-line message software is supplied in busex.c and consists of the following processes:

do_bus_ex(). This is the foreground process; it accepts data entered from the keyboard by the user (sent from the keyboard interrupt handler to the APPLICATION queue). Do_bus_ex() also handles the message response codes returned from the bus network.

Keyboard data is received a byte at a time from KEYIH. All bytes (except CTRL-C and Enter) are stored in the array buffer and echoed to the screen (using disp1ay()). Receipt of an Enter causes the buffer to be sent to the Busnet queue, while a CTRL-C results in the bus network software (on this station) terminating.

Once the bus network has completed the transmission, a message code is returned to the foreground process indicating the result of the transmission (MESSAGE_SENT, COLLISION_LIMIT, or BAD_MSG_SIZE).

do_display(). Messages sent to the BACKGROUND_1 queue are read by the do_display() process and displayed in angle brackets on line 10. The message displayed also includes the destination and source addresses, permitting the verification of any message received.

11.4
Using the Commkit Bus Network

The Commkit distribution diskette is supplied with both source and executable versions of the bus messaging software.

11.4.1
Creating the Bus Network Software

The Commkit makefile contains instructions to the make utility for the creation of the messaging software, busex.exe. The high-level software is linked with buslow.obj, commkit.obj, crcsr.obj, and srcalls.obj

using the `make` utility. For example, to create bus `ex.exe`, type

```
C:\> make busex.exe
```

The Turbo C compiler includes `general.h, ascii.h`, and `devices.h` with `busex.c` to create `busex.obj. Busex.obj` is then linked with `buslow.obj, commkit.obj, crc.obj`, and `srcalls.obj`, producing busex.exe.

11.4.2
Running the Bus Network Software

Once a compiled version of the bus network emulator is available, it can be run by typing busex after the DOS prompt. The line speed (50, 300, 1200, 2400, 4800, or 9600 bps) and the station identifier (a single character) must be entered on the same line. For example, to run the bus messaging software at 9600 bps and to set the station's identifier to S, type the following:

```
C:\> busex 9600 S
```

If the format above is followed and the line speed is acceptable, the screen is cleared and message transfer can commence. Error messages are issued if the format is incorrect or the line speed is not recognized.

Since the software is running at a low level, it is necessary to type the destination and source address as part of any message. For example, if the user at station S is to send the message "How are things at your end?" to station L, the following is typed (note that there is no prompt: ENTER denotes end-of-message):

```
LSHow are things at your end?
```

Packets with destination addresses matching the address of the station are displayed on line 10. For example, at station L the above packet is displayed as follows:

```
LSHow are things at your end?
```

To understand what happens during a transmission or a collision, it is best to run the line analyzer on a PC connecting two bus network PCs. Collisions can be caused by typing messages on both PCs (the destinations are not important) and pressing Enter simultaneously on both. Although the two Enter keys may be pressed at the same time and the same algorithm is running on each PC, a collision may not occur, since there is no guarantee that the clocks are synchronized. If collisions seem impossible to achieve, rerun `busex.exe` at a lower speed.

11.5 Summary

This chapter examined bus local area networks. Bus networks are named after their topology, in which all stations share a single communication channel. There are essentially two problems that designers of bus networks must overcome: how to control access to the network and how to resolve collisions. Access control is handled by requiring each station to listen to (or sense) the line prior to transmitting. Collision resolution is handled first by having the stations determine that a collision has occurred and then cease their transmissions; and second by minimizing the chances of a collision happening again. A common solution is have each station back off from transmitting for potentially longer and longer times by picking random waiting periods. Many bus networks (such as the Ethernet and IEEE 802.3) implement such an algorithm, known as truncated binary exponential backoff. Much of the communication handling is now performed by VLSI chips.

The Ethernet and IEEE 802.3 networks are often referred to as CSMA/CD (Carrier Sensed Multiple Access with Collision Detection). That is, the network is sensed before transmission and if a collision is detected, transmission is halted.

11.6 Exercises

All of the exercises relating to the bus network software assume access to at least three PCs, each with two serial ports.

In exercises where the bus network software is to be modified, try testing the modified software against another modified copy.

1. (B) Set up a 3-station bus network and run the bus network software on the two outermost stations and the line analyzer on the middle station. Assign each station on the bus a unique address (such as A and B). Send packets between the two stations (remember that at this level both the destination and source addresses must be embedded in the packet). Examine the bus network protocol using the line analyzer by transmitting a packet from one station to the other. Identify the various fields within the packet. (The *STX* character is displayed as , while the *ETX* character is .)

What happens if a packet is sent to station C (i.e., a nonexistent station) from either station A or station B? Does the packet appear on the line analyzer? If so, why?

2. (B) Set up a 3-station bus network and run the bus network software on two adjacent stations and the line analyzer on the third station. Assign each station on the bus a unique address (such as A and B). Now what happens if a packet is sent to station C from either of the stations? Does the packet appear on the line analyzer? If so, why?

3. (B) Configure the network with the line analyzer running between two bus network stations (as in exercise 1). Type messages on both stations and attempt to transmit them simultaneously (this takes a bit of practice and is usually best achieved by having one person press the Enter key on both stations).

If the packets are sent simultaneously, a collision should occur and be displayed on the line analyzer. (If, no matter how many attempts are made, a collision never occurs, rerun `busex.exe` at a lower line speed.) What happens during the collision? Which station "wins" and is allowed to transmit?

Start up both stations with the same network address; what happens once a collision has been detected? Given the backoff algorithm used in the Commkit bus network, what combination of addresses will result in equal backoff time values?

4. (C) As discussed in this chapter, a broadcast communication is one in which a single packet is delivered to all possible stations on a network. In a bus network, all transmissions are naturally broadcast since all stations must receive at least the destination address before deciding whether to keep the packet.

The Commkit bus network is also a broadcast network; however, since a broadcast address is not supported by the network, broadcast messages cannot be sent. Add broadcast addressing to the Commkit bus network. Use `0xF0` as the broadcast address (`0xF0` can be generated by typing ALT P, which echoes as ≡).

Note that sending a message with a broadcast address is relatively simple (the unicast address is replaced by the broadcast address during the transmission). However, it is during reception that the check must be made.

Test the resulting software on a 3-station network. First, try a broadcast transmission by sending broadcast packets and watching the traffic with a line analyzer. Then test broadcast reception by setting up one (then two) stations to receive and send a broadcast packet from the third. Are collisions still supported? Does unicast transmission still work? Does the station sending the broadcast packet also receive a copy of the packet?

5. (C) A multicast address is one that is *shared* by a number of stations on the network (from none to all possible stations). Implement multicast addressing on the Commkit bus network. Allow a station to belong to a maximum of 10 multicast addresses.

Since multicast addresses can be turned on and off at random, it will be necessary to make some relatively major changes to the bus emulation software because, at present, the only reason for communicating with the bus software is for the transmission of a message. Consider setting up a new data structure (possibly a union), in which the first byte indicates to the bus network the action that is to be taken (i.e., whether the message

contains data to be sent or a new multicast address to be added to the list).

In the Ethernet, a multicast address is distinguished from a unicast address by the most significant bit being set to 1. Use a similar scheme for this question: give multicast addresses a value between `0xE0` (ALT ') and `0xEF` (ALT O); the multicast addresses are echoed as the symbols *a*, *b*, Γ, *p*, Σ, *s*, *m*, *t*, Φ,Θ, Ω, *d*, ∞, ϕ, ∈, and ∩. Remember, the broadcast address, `0xF0` (≡), is accepted by all stations.

6. (C) Another type of network analysis tool that is available to many network managers is the network analyzer, which is similar to a line analyzer in that data can be captured and displayed; however, most network analyzers deal with packets rather than individual bytes. A typical network analyzer will allow a network manager to capture packets as they are sent across the bus network, regardless of the destination. Some allow the user to selectively capture packets, for example, by monitoring a particular source or destination address or even scanning for certain byte values within the packet. In order to capture data destined for any station on the network, the network analyzer operates in promiscuous mode. Write a network analyzer for the Commkit bus network. As packets are received, they are displayed on the screen (use `printf()` to obtain scrolling). Add enhancements to the analyzer that allow the user to request a specific source or destination address to be searched for.

7. (C) At present, the Commkit bus network has its own packet structure, unlike either the Ethernet or the IEEE 802.3 packet structure. Modify the Commkit bus to support a packet structure similar to that of the Ethernet by adding a 1-byte Type field to the Commkit bus packet. The new version of the software should route the packet to the process associated with the value supplied in the Type field.

Test the software by having two processes on a station, each associated with a different Type field. Messages destined for the process associated with Type field value `0x45` should be displayed in normal video, while messages sent with Type field value of `0xF4` should be displayed in inverse video. Remember, the destination address takes precedence over the Type field (that is, the Type field should only be considered if the destination address is the same as that of the station).

Avoid putting the Type field check directly into the bus network driver. Instead, send the message to a background process that forwards the message to the intended destination process. Consider a number of tests, such as sending a message with a Type field value that is not supported on the station to which the message is sent.

8. (C) The Ethernet uses 48-bit addresses to identify a destination host, whereas Commkit uses a single 8-bit address (since it was assumed that no one would attach more than half a dozen stations to the network). Modify the Commkit bus emulator so that it supports 48-bit addressing.

Consider techniques whereby an address can be checked on the fly, a byte at a time.

9. (C) Rewrite the bus network's backoff algorithm so that it uses the Ethernet's truncated binary exponential backoff algorithm. In the original implementation of `buslow.c`, the Ethernet algorithm was used and subsequently abandoned. Why?

10. (C) Implement the point-to-point telephone utility atop the bus network. Each byte entered will have to be sent in its own packet. Test the software against another modified copy.

11. (B) Develop a testing methodology that can be used to cause each of the bus network error messages to be produced.

12. (B) Show why the minimum Commkit bus network packet size is $2 \times N$. where N is the number of stations.

13. (A) A number of papers have been written that claim that unicast and broadcast addresses are simply special cases of multicast addresses. Explain whether and how this might be true.

12
Ring Local Area Networks

12.1
Introduction

The CSMA/CD bus, for all its benefits, is by no means the only or necessarily the best design for a local area network. In the worst case, it is possible that with the bus LAN a station may never be able to transmit its data. Consider the following scenarios:

Two stations attempt to transmit simultaneously and a collision occurs. By some coincidence, both stations generate the same backoff values, causing continual collisions. Eventually, the collision limit is reached, and the transmission of at least one of the packets is aborted.

A heavily loaded network consists of several hundred stations all trying to transmit large volumes of data simultaneously. Initially, all transmissions will result in collisions; however, as backoff delays begin to grow longer, some of the transmissions will be successful. Stations that constantly collide might never get an opportunity to send their data.

These situations are extremely unlikely on a lightly loaded bus network, but they have the potential for seriously degrading a heavily loaded network. The source of the problem can be traced to the nondeterministic nature of the bus: access to the bus does not guarantee that a station can transmit.

12.2
The Token Bus

To ensure that every station with data to transmit on a bus network has a guaranteed opportunity to transmit implies some form of deterministic network access, which requires a certain degree of control that the bus network by itself does not offer. One possible approach is to have a station dedicated to polling all other stations for transmissions, thereby ensuring fair access.

However, in Chapter 11 it was shown that centralized polling has its limitations, an important one being how to allow access to the network should the polling station fail.

A variation on polling is to allow each station to transmit in turn, but instead of a single, centralized polling station, permission to access the network is passed *between* the stations themselves in the form of a token. Stations no longer compete for the network; instead, they wait for the token, at which point they can transmit. Once transmission has completed, the transmitting station is responsible for forwarding the token to another station to access the network. This cycle continues indefinitely.

At any time, the packet on the network contains either data or a token. Since stations may not transmit unless they have the token, collisions should be rare or nonexistent.

Guaranteeing deterministic network access using the token method means that after any station has had its opportunity to transmit, it must wait until all other stations have been given the same opportunity. For example, in a network of four stations (A, B, C, and D), once A has transmitted, it cannot transmit again until B, C, and D have been given permission to transmit. The question is: what type of mechanism is required to enforce this ordering?

As an analogy, consider the path of a memo through an office. The memo originates from the secretary, who passes it to the person named first on the list. That person passes it to the second, the second to the third, and so on, until it reaches the last person, who is responsible for returning it to the secretary. The memo's path could look like that shown in Figure 12.1.

The last person on the list is expected to return the memo to the secretary. Similarly, in the network of four stations the last station with the token must forward it to the first station. The path taken by the memo (or the token) is similar to that required to traverse a circularly linked list.

The bus network is not a circular list (it is an unrooted tree); however, it is possible, through software, to implement a circular linked list atop the bus. Quite simply, every station is followed by a unique "downstream" station. Figure 12.2 is an example of how a circularly linked list could be organized on top of a bus network; Table 12.1 shows the linking of the stations in the figure.

If each station "knows" the station that follows it, forwarding the token is a simple operation, as illustrated by the following algorithm:

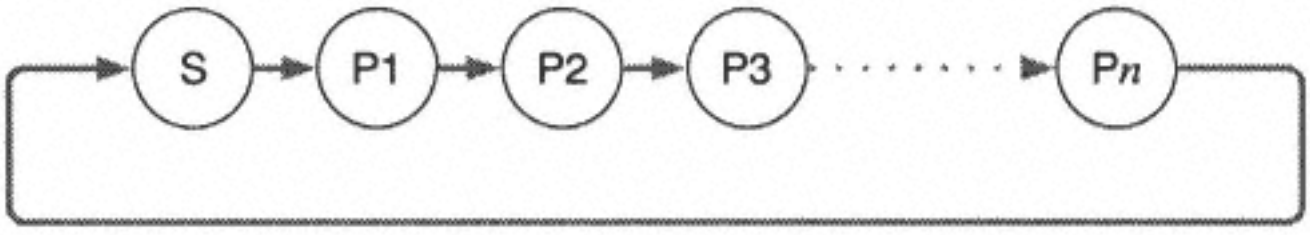

Figure 12.1
Path of a memo through an office.

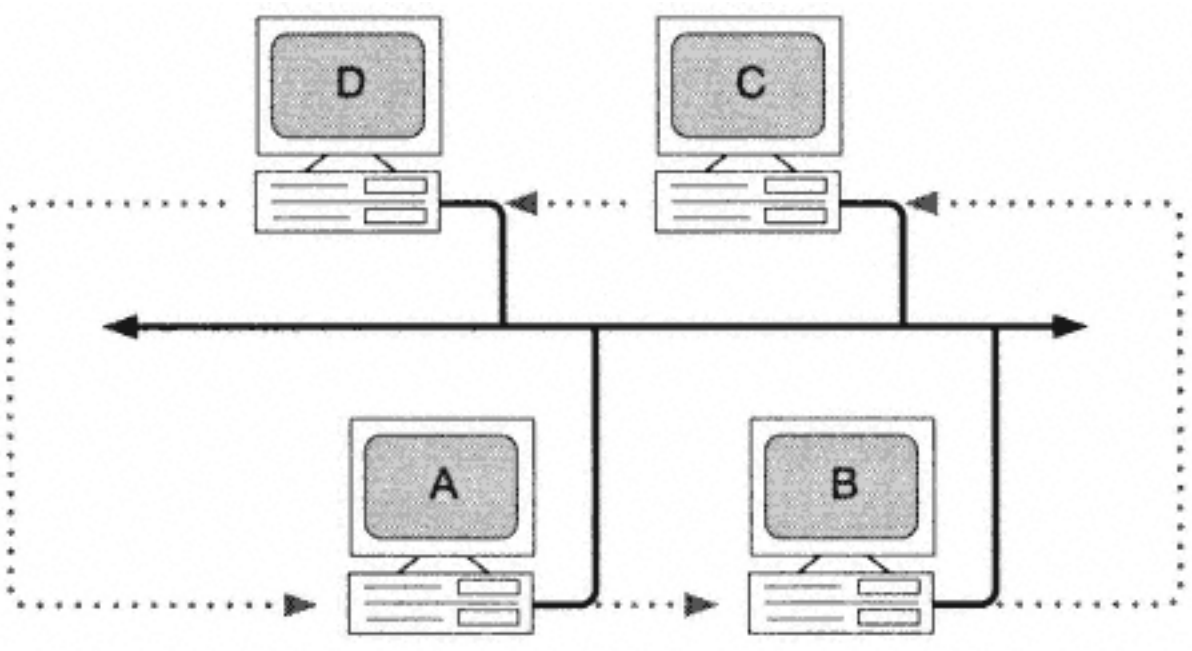

Figure 12.2
Imposing a linked list on a bus network.

1. Wait for the permission-to-transmit token from the "upstream" station.

2. Upon receipt of the token, transmit any pending messages.

3. Forward (i.e., transmit) the token to the "downstream" station.

4. Repeat from step 1.

Access to the network is now deterministica station is guaranteed access once all the other stations have had an opportunity to transmit. Note that the underlying bus network remains unchanged; it is still a broadcast network (any packet sent on the network is received by all other stations and discarded by them if its destination address doesn't match their address). The bus now supports two types of packet:

Data packets. Packets containing information sent from one station to any other station(s). A data packet is simply a bus network packet carrying data.

Token packets. Packets containing an indication that the station specified in the destination address may access the network. A token packet is a bus network packet carrying token information from one station to another.

The use of the token to permit network access ensures that there is an upper limit on the time a station must wait before it can transmit. In the worst case (i.e., all stations transmit a data packet when the token is received), a station will have to wait a total time equivalent to

(*Number of stations*1) × (*Time taken to send a data packet*)

Table 12.1 Station Linkage for Figure 12.2.

Station		Downstream
A	→	B
B	→	C
C	→	D
D	→	A

12.2.1
IEEE 802.4: Token Bus

The push for a token bus standard came, in part, from some of the issues raised in the introduction to this chapter: the bus is nondeterministic, meaning that stations on the network may be required to wait unduly long periods before being able to transmit a frame. The 802 standard committee recognized this and set about developing another standard: IEEE Standard 802.4, Token Bus. Two main proponents of the 802.4 standard were General Motors and Boeing (both interested in factory and office automation), which were involved in the development of MAP (Manufacturing Automation Protocol) and TOP (Technical and Office Protocols), respectively.

The 802.4 protocol is *not* built atop the 802.3 CSMA/CD MAC layers. Instead, it defines its own physical and MAC layers. The physical layer uses a broadband backbone coaxial cable and provides three data rates (at separate frequencies): 1, 5, and 10 Mbps. The separate rates are intended for equipment that supports data, voice, or video. *Physically*. the 802.4 token bus is a broadcast network; *logically*. the stations on the bus are connected as part of a circularly linked list.

The 802.4 frame format is shown in Figure 12.3. The fields in the 802.4 frame are defined as follows:

Preamble. A clock synchronizing sequence; its minimum size is 1 byte.

Start Delimiter and End Delimiter. Marks the start-of-frame (or end-of-frame) by encoding electrically invalid Manchester bit patterns into the byte. The two delimiters are not the same, although both are one byte long.

Frame Control. Denotes the use of the current frame, either data or control. When signaling a data frame, the Frame Control field indicates the frame's priority (see below) and whether an acknowledgment is required by the transmitting station.

Destination and Source Address. The 802.4 frame supports an addressing scheme identical to that of the 802.3 CSMA/CD bus. Both 16- and 48-bit addresses are supported, although not simultaneously in the same network.

Data. The message field, up to 8182 bytes long (when using 16-bit addresses) or 8174 bytes long (when using 48-bit addresses).

Checksum. Uses the 802.3 frame check sequence algorithm and generator polynomial; it includes the entire frame (between frame delimiters).

The 802.4 MAC layer is responsible for two main tasks: data transfer and network maintenance:

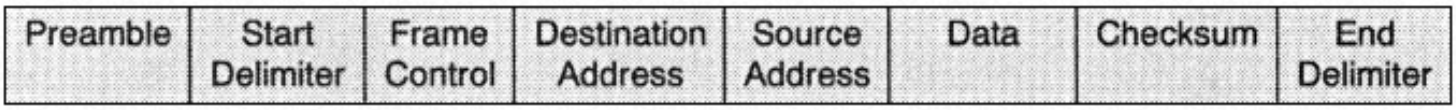

Figure 12.3
The 802.4 frame format.

Data Transfer. When a token arrives at a station, the station has a certain amount of time available in which to transmit messages in 802.4 frames; this time is divided between various message queues. The token bus defines four levels of message priority (0, 2, 4, and 6, with 0 being the lowest and 6 the highest), and each level is associated with its own queue. Messages are transmitted from the highest-priority queues first. Once the messages in a queue have been sent or the queue timer has expired, or if there was nothing in the queue to send, the next lower-priority queue is given permission to start its transmissions. If the station has nothing to transmit, has finished transmitting, or its time limit has expired, the token is forwarded to the station's successor, a control frame with the control frame field set to `TOKEN`.

The order of token passing is strictly defined: from high to low addresses. Each station maintains the address of its predecessor and its successor.

Network Maintenance. Network maintenance covers the addition of new stations to the network, the removal of stations from the network, and the initialization of the network. It is governed by the stations themselves, through control frames. The conditions that 802.4 can handle include

Lost token. A token bus cannot function if there is no token on the network. The `CLAIM_TOKEN` control frame code is used when a station determines that the token is missing and it is necessary to place a new token onto the network; for example:

1. When a station is first brought on-line, it monitors the network for traffic; if none is detected within a fixed period of time, the station assumes that it is the first station on the network. However, before putting a token onto the network, the station first broadcasts a `CLAIM_TOKEN` frame. If no responses are received, it sets up a network of one (itself) and sends frames addressed to itself (thus generating network traffic).

2. All stations contain timers monitoring the network for the token; if a token is not detected within a certain period and a station's timer expires, the station broadcasts a `CLAIM_TOKEN` frame. If collisions occur during the bid for the token, various timers come into play, allowing one station to eventually win. At this point, a new token is placed on the network.

Adding stations. The physical presence of a station on the network does not necessarily mean that it will receive a token; the station must "join" the network and become the successor to one station and the predecessor of another.

The `SOLICIT_SUCCESSOR_1` control frame is sent periodically by any station that has the token, thereby allowing other stations to participate in communications. In order that the strict highest-to-lowest sequencing of stations is maintained, the frame contains the sender's and successor's addresses; only stations waiting to join the network with addresses in that

range are allowed to bid. If no bids are received within a certain period, the station with the token proceeds with its data transfer. If a bid is received from one station, it becomes the successor of that station. Should several stations bid simultaneously, their bids will collide, requiring the use of a `RESOLVE_CONTENTION` frame to indicate that the bids have failed. The station holding the token then attempts to add a new station by halving the range of possible addresses (i.e., the mid-point between its address and the address of its successor). This cycle continues until a successor is found.

Stations leaving the network. When station Q is about to leave the network, it sends a `SET_SUCCESSOR` frame containing the address of its successor (say, R) to its predecessor (P). P makes its new successor R, and Q is now removed from the network.

Missing stations. When the station currently holding the token passes it to its successor, it monitors the bus for subsequent traffic. Ideally, there will be traffic generated by the successor; however, if nothing is heard, the original token holder sends the token again. If no response is heard the second time, the token holder assumes the worst and broadcasts a `WHO_FOLLOWS` frame containing the address of the missing station. The missing station's successor responds with a `SET_SUCCESSOR` frame, and the token holder updates its successor address and forwards the token to that station. If no responses are forthcoming to `WHO_FOLLOWS`, the token holder broadcasts a `SOLICIT_SUCCESSOR_2` frame. Stations wanting to join the network can bid to become the successor.

One of the reasons for the 802.4 topology is that it lends itself well to automated-factory assembly lines using robots or other computer-controlled equipment. The cabling connects the various pieces of equipment.

By completely changing the topology of the network (that is, by abandoning the bus structure), it is possible to achieve results similar to those of the token bus, with a somewhat less complicated protocol and potentially higher throughput. This topology is the ring.

12.3 Ring Networks

The topology of a ring network is as the name suggests: a ring interconnecting all stations, illustrated in Figure 12.4. Accessing a ring is similar to accessing a token bus: the station must wait until permission is granted to transmit. However, there are several notable differences between the ring and the token bus:

Physical topology. The token bus is a deterministic ring built atop a nondeterministic bus network; all transmissions are broadcast. As it is passed from station to station, the token follows no particular direction; the physical location of a station has no bearing upon when it receives the token.

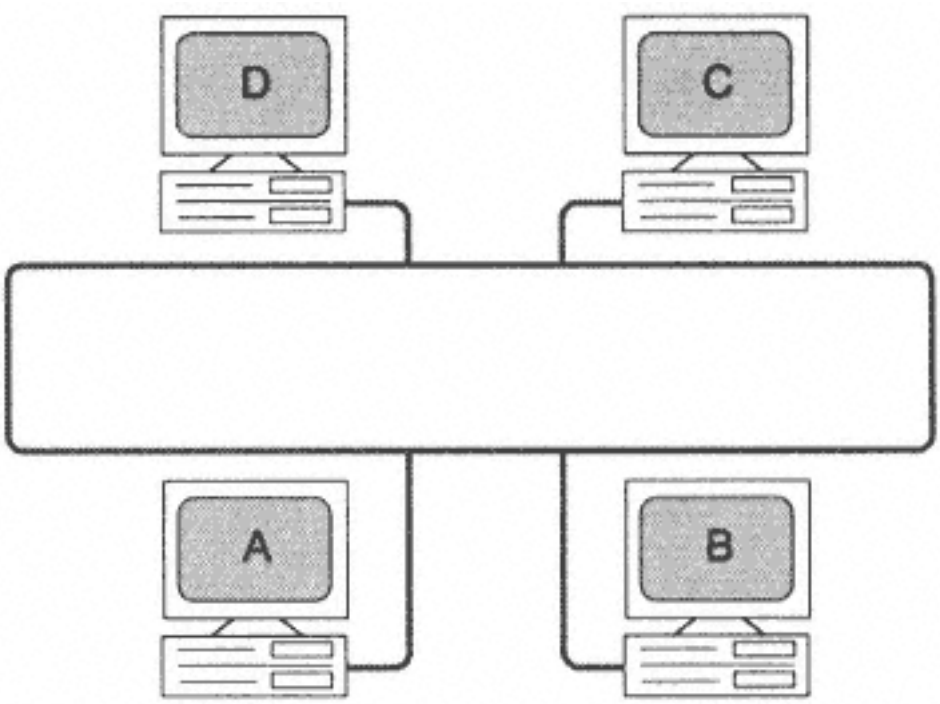

Figure 12.4
The basic structure of a ring network.

A ring network consists of a number of stations connected in a closed loop. In most rings, packets flow in one direction onlyout from the transmitting station, past all the other stations, and back to the transmitting station, which is responsible for removing the packet.

Successor identification. The successor station in a token bus must be explicitly identified in the token packet, whereas in a ring the successor is the next station "downstream" from the current one.

Reliability. Most token bus and ring networks are designed to continue functioning should a station fail. However, if the channel fails, the results can be quite different. In theory at least, the bus can continue functioning in a degraded fashion as two separate buses.

A channel failure on a ring can result in the total cessation of communications, since all information flows in one direction around the ring. Recent developments in ring architecture use *two* bidirectional channels: should one fail, packets can be routed through the other.

There are, broadly speaking, three types of ring network algorithm: the token ring, the slotted ring, and register insertion.

12.3.1
The Token Ring

Accessing a token ring is similar to accessing a token bus: stations must wait until they receive permission before transmitting a message. The ring itself is in one of two states: free or busy, depending upon the value of the token in a variable-length frame that is continuously circulating around it. Frames are made up of bits, and at any moment there are a limited number of bits on the ring. The station hardware can sense and change the value of a single bit as it passes the station.

Token ring frames consist of a start-of-frame indication followed by a control byte that contains a single bitthe tokenindicating whether the

ring is free or busy. The free token frame ends after the control byte with an end-of-frame indication.

A station detecting a free token (i.e., the token bit within the control byte is sensed by the ring physical-layer hardware as free) is able to transmit a message. This is done by the ring hardware first setting the token bit within the frame to busy (the bit then continues on). At the end of the control byte, the destination and source addresses are sent. Then each bit in the message is transmitted, followed by a CRC and the end-of-frame indication. The frame circulates around the ring (typically only a few bits are on the ring at any one time), passing each station.

Upon detection of a busy token, each station takes a copy of the frame, assembling it out of the sequence of incoming bits. When the end-of-frame is detected, the receiving station can examine the destination address of the packet to determine if it should be kept or discarded.

The transmitting station is responsible for removing the bits from the ring (since only a few bits are on the ring at any one moment, they will interfere with the transmission of the remaining bits if they aren't removed). When the transmitting station has completed its transmission, a new free token (consisting of the start-of-frame, control byte, and end-of-frame indication) must be put on the ring.

Figures 12.5, 12.6, 12.7, and 12.8 illustrate the cycle of events as station D transmits a message to station B. First, in Figure 12.5 station D waits for a free token. Once the free token is detected, station D sets the token to Busy and starts to transmit the remainder of the frame (the destination address, the source address, the message, and any checksum). Note that during part of the transmission, some of the bits of the free token are still being removed from the ring by station D (Figure 12.6).

The packet circulates past the stations on the network; since the token bit indicates Busy, each station assembles the remainder of the frame. When

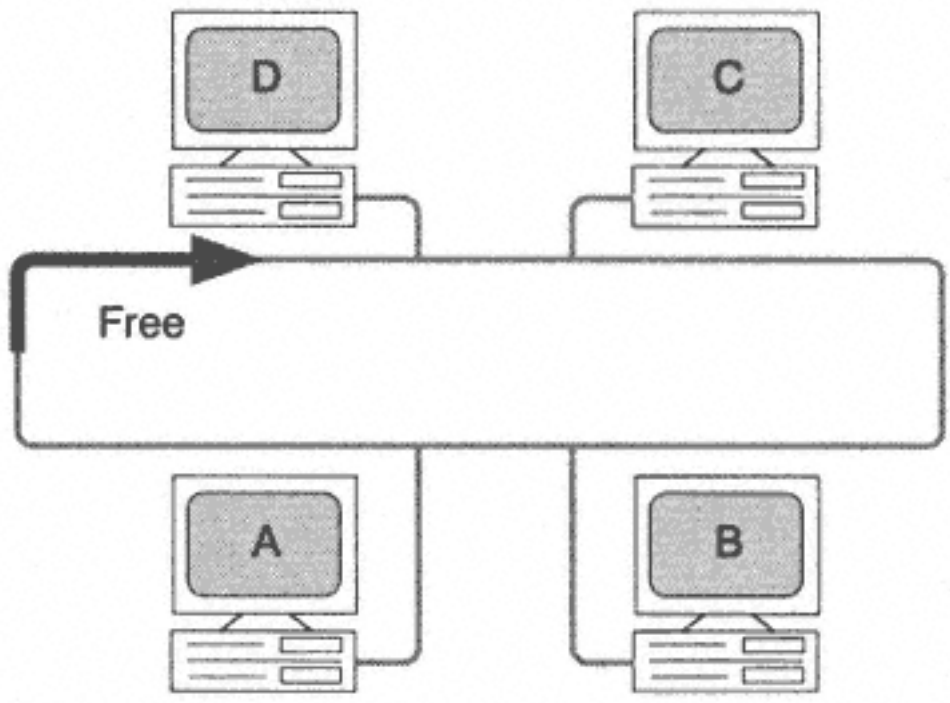

Figure 12.5
Station D awaits the Free token.

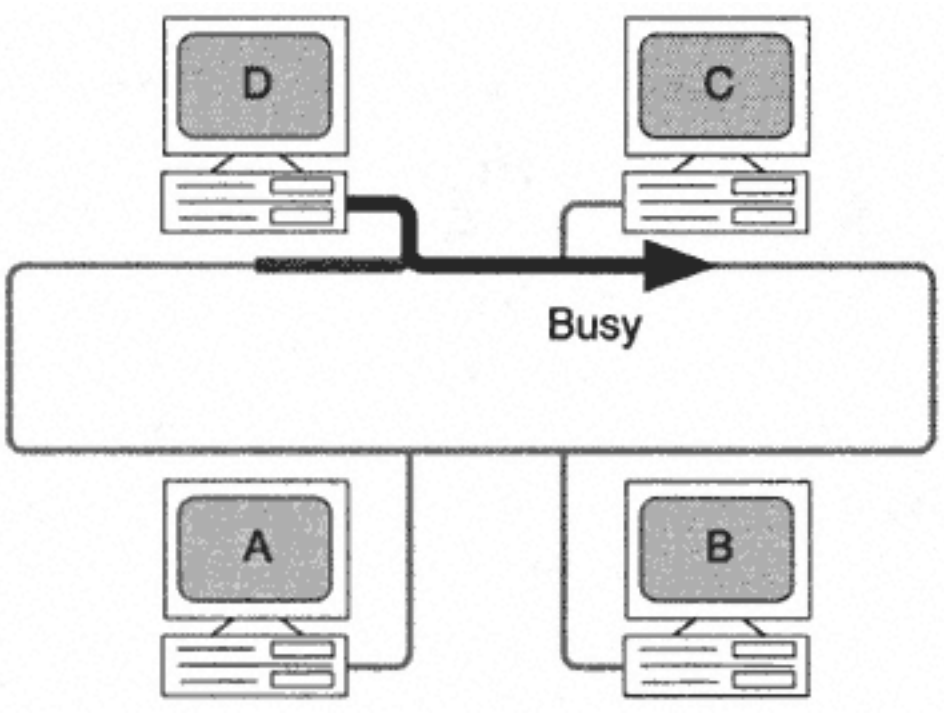

Figure 12.6
Station d changes the token from Free to Busy.

the destination address has been assembled, each station compares it with its own: if the destination and the station's address are the same, the station continues to accept the bits making up the frame. Upon detection of end-of-frame, the CRC is checked and the message is forwarded to the station. In this example, only station B takes a copy of the frame since the destination address indicates B (Figure 12.7).

If the number of bits in the packet exceeds the number that can circulate at any moment on the ring, the transmitting station (D in this example) removes the bits while it is transmitting. It continues to transmit until the end-of-packet is sent, at which point it can apply one of the following algorithms:

Remove the entire packet from the network before reissuing the free token. This algorithm is known as single-frame operation, and it reduces the throughput of the network since there are periods in which the network is devoid of any frame.

Issue a new free token as soon as the busy token is removed from the ring

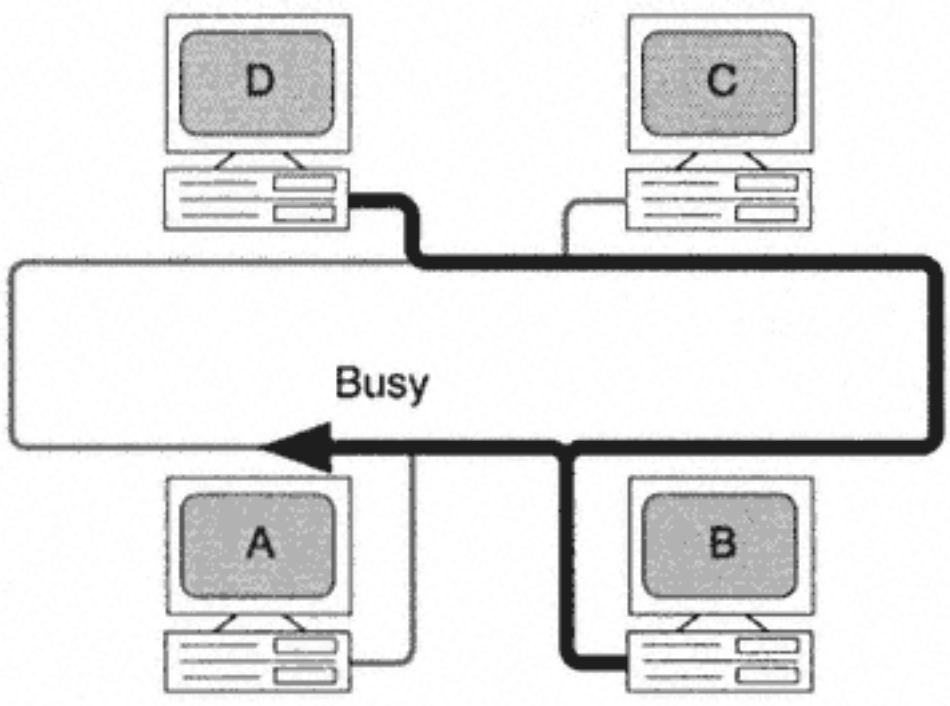

Figure 12.7
Station B keeps a copy of the frame.

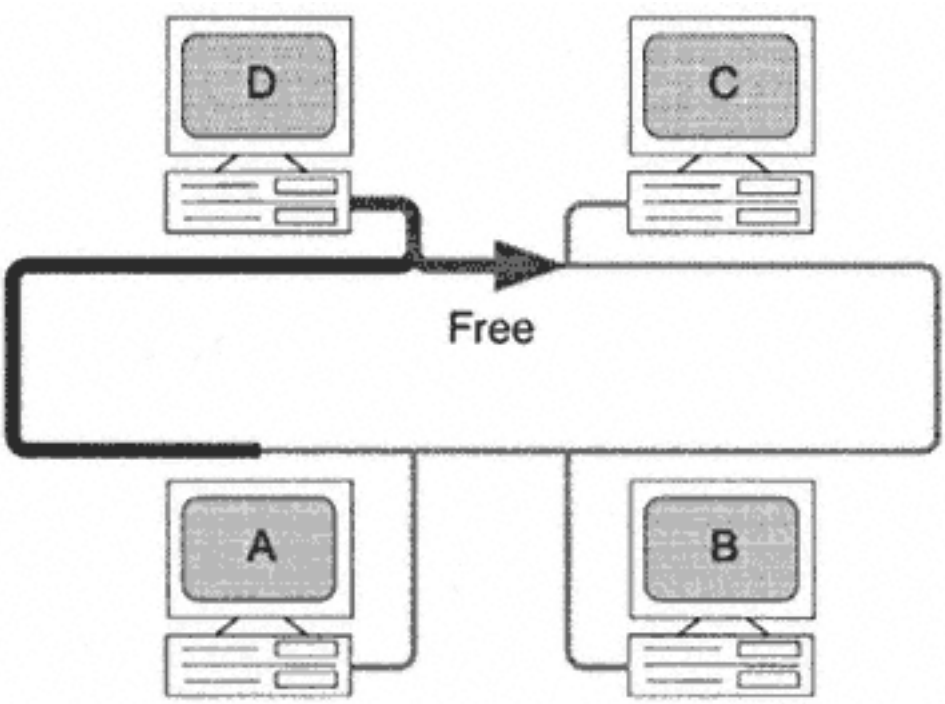

Figure 12.8
Station D is responsible for freeing the token.

and the end-of-frame has been sent. This algorithm, known as single-token operation, can be used when the number of bits in the frame is less than the number that can exist on the ring at any time.

Start to transmit a free token as soon as the end-of-frame is sent, which implies that several tokens can exist on the ring at any moment. This is known as multiple-token operation, and it ensures that the amount of network idle time is kept to a minimum.

In this example, station D issues a free token *before* the original frame is completely removed from the network (i.e., the network supports single-token operation) (Figure 12.8).

The above example, which requires the transmitting station to remove the frame as it is being transmitted, is known as source removal. The alternative is to have the destination station remove the frame as it is received (known as destination removal). Although destination removal suggests that greater throughput can be achieved, it suffers a number of drawbacks, such as the following:

The deterministic nature of the token ring can be lost. That is, instead of ensuring that each station will be able to transmit within a certain number of frame cycles, one station can continue to transmit without giving up the ring. For example, if station A transmits frames to station B, and B issues a free token, then A will receive a free token before any downstream stations (D, C, or B) have had a chance to transmit. Station A can therefore monopolize the network.

Support for broadcast and multicast communications is cumbersome. For example, if the first station to receive a broadcast frame releases it as a free token, no other stations receive a copy of it. This problem can be overcome by requiring the transmitting station to perform source removal on broadcast and multicast frames only.

A simpler solution is to use source removal for all transmissions, thus eliminating these special cases.

Ring Errors

The objective of a ring network is to ensure that all stations have equal ring access. This is achieved with some form of circulating token. Should anything happen to the token to stop it from circulating, network access may become impossible. Two error situations can stop network access entirely:

Missing token. If the token is removed from the network or becomes so corrupted (through noise) that it is impossible to recognize, all stations will be left waiting for a free token that never arrives.

In addition to being lost because of noise, the token can go missing if a station with a message to transmit removes the free token from the ring and *crashes* before putting the busy token on the ring. Also, the token can disappear if a frame has been sent and the transmitting station crashes after removing the busy token but before putting the free token back.

Circulating busy token. The token remains in the busy state, either by being corrupted (i.e., changed from free to busy, typically because of noise) or by a transmitting station crashing after setting the token to busy and so being unable to put a free token back on the network.

Should either one of these situations occur, the network will cease to function. To avoid the problem, some form of centralized control is needed. Most rings employ a station to *monitor* the network for either of these conditions.

The Monitor Station

The solution adopted by most ring network designers is to assign one station the task of monitoring the ring to ensure that if the network enters an error state (such as a missing or a circulating busy token), the error can be corrected and a new token put on the ring. This station is known as the monitor station; its relationship to the ring and the other stations on it is shown in Figure 12.9.

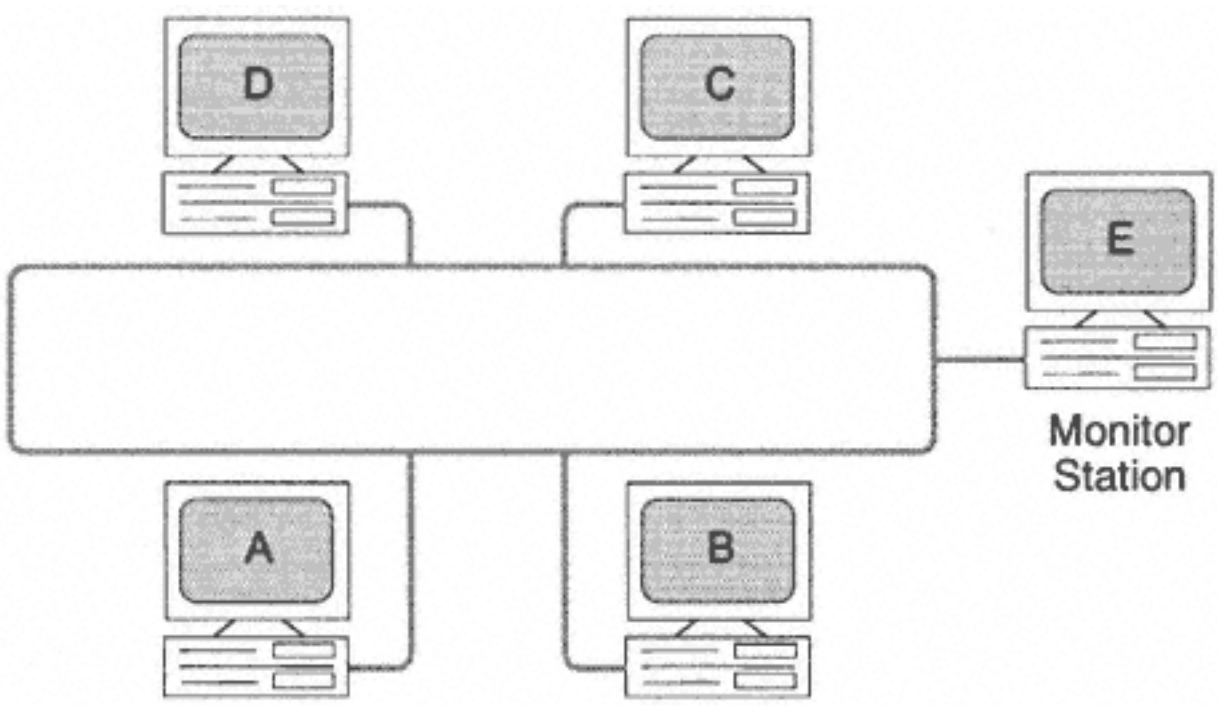

Figure 12.9
The ring monitor.

Depending upon the network, the monitor can be dedicated to monitoring or it can perform monitoring in addition to the transmission and reception of frames.

Regardless of how the monitor is implemented, it must be able to handle the two error situations described above. The monitor can detect the missing token situation with little difficulty, as the network is always being monitored for traffic. Tokens (busy or free) are expected at regular intervals, so that if one does not arrive within the designated period, the monitor can assume the worst and issue a new free token frame.

Detecting the circulating busy token is somewhat more difficult, requiring the monitor to determine whether the busy token has just been set to busy or is still busy from a previous cycle around the ring. The approach taken by ring networks is to add a bit to the control field set by the monitor each time a busy token bit is detected. This bit, often referred to as the monitor passed bit, signals that a busy token frame was detected by the monitor. The monitor is responsible for setting the monitor passed bit when it receives a frame with the token bit set to Busy and the monitor passed bit cleared; this is shown in Figure 12.10 (MP denotes "Monitor Passed"). The transmitting station is responsible for issuing a free token frame with the monitor passed bit cleared.

There are two conditions that can arise when the monitor detects the busy token: the monitor passed bit is either set or cleared:

Set. The station transmitting the packet did *not* remove the frame from the network. This error condition requires the monitor to remove all bits from the ring and put a new free token frame on it.

Cleared. This is the first time the packet has passed the monitor, which then sets the monitor passed bit.

Rings can be of varying sizes (i.e., total length or number of stations), and the size can determine the number of bits on the ring at any one time. In many rings, the monitor is also responsible for maintaining an elastic buffer, which, acting like a queue, regulates the number of bits on the ring. For example, a free token may require 24 bits; if the ring has space for only 8, the elastic buffer must be holding a queue of 16.

Figure 12.10
The monitor station indicates "monitor passed" if the token is Busy.

Other Ring Errors

Rings that allow at most a single token (free or busy) on the ring at any time (i.e., they operate in single-frame or single-token mode) must be protected from conditions in which there is *more* than one token on the ring. The duplicate token situation can be caused by noise or some other error and can, for example, result in two stations attempting to transmit simultaneously.

The solution to the duplicate token situation requires each transmitting station to examine the source address of the frame *before* removing the frame from the ring. If the source address is not the same as the station's address, a duplicate token situation has occurred. Once the situation is detected, a number of possible algorithms can handle it:

If all stations simply remove frames that do not have the correct source address and do not place a new free token on the ring, the monitor station will eventually detect a lost token situation and reissue a new free token frame.

The previous solution results in idle periods in which no busy frames can be transmitted until the monitor puts a free token on the ring. An alternative solution is for the stations to examine the frame's source address and for the one with the *lower* address value to stop its transmission. (To do this, each station maintains a queue of the incoming bits.) The station with the higher address value continues to transmit its frame.

The station with the lower address value forwards the bits already in its queue (starting with a start-of-frame indication). Subsequent bits are taken off the ring and stored in the queue for transmission, ensuring that the bits are received in the correct order by the transmitting station.

A second error that must be considered in any ring network that uses a monitor to control network access is what happens if the monitor fails. Should this occur and an error situation such as the missing token arise, network access will cease entirely.

Recovery from a monitor failure requires two steps: first, the loss of the monitor must be detected; second, once detected the loss must be recovered from. Detection can be carried out in a number of ways, for example:

The monitor can be required to periodically send a "monitor alive" message to all stations on the network; if the message is not detected within an agreed-upon time, the monitor has failed.

Since access to the ring is deterministic, each station knows that within a certain maximum time, a free token should be available; if the token is not found, the monitor has failed.

All stations can monitor the ring for traffic; if nothing is detected within a given period, the monitor has failed.

Once the error has been detected, the remaining stations must determine which of them is to become the new monitor (assuming that a station can

become a new monitor); otherwise, the error must be signaled to the person managing the network and the monitor repaired manually.

In a typical recovery algorithm, each station that has detected the loss of the monitor announces this discovery (by means of a frame). As the frame circulates around the ring, stations that have discovered and announced the loss of the monitor examine the frame's source address. If the address is *less* than the station's address, the frame is discarded; that is, the extra frame is handled as a duplicate token situation. The station that receives a packet with its own address as the source address becomes the new monitor.

It issues a new free token on the ring, and the network resumes activity.

Connecting to the Ring

In most ring networks, the ring does not pass physically through the stations connected to it. If it did, the failure of a station would cause the entire ring to cease functioning. Instead, each station is connected to the ring through a bypass relay (shown in Figure 12.11). As long as the station is electrically active, the bypass relay remains open, allowing the station to receive individual bits. Should the station fail, the bypass relay closes and the bits avoid the station, flowing past it.

Examples

The previous section's description can be applied to most token rings. However, a number of enhancements and extensions found in current ring implementations are listed below.

The IBM Token Ring functions for the most part like the generic token ring described earlier, but it has several notable features, described in the paragraphs that follow.

The busy token frame format is illustrated in Figure 12.12 (all sizes are in bytes). The frame delimiters, DEL, shown in Table 12.2, are bytes containing invalid Manchester bit encodings and different bit values to distinguish between the starting and ending delimiters (X denotes an invalid bit pattern).

The Control field, CTRL, consists of two bytes. The first is common to both the free and busy token frames and consists of the fields shown in

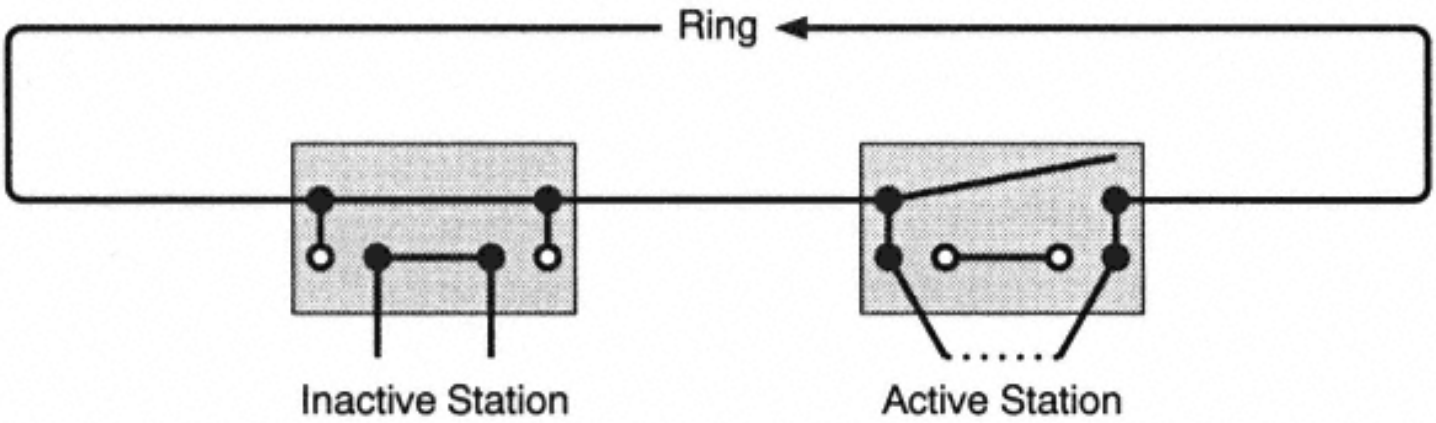

Figure 12.11
The bypass relay.

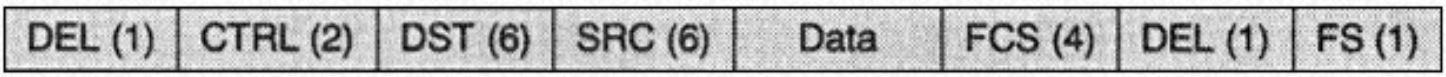

Figure 12.12
The busy token frame format for the IBM Token Ring.

Table 12.2 Encoding of DEL for the IBM Token Ring Frame.

Starting delimiter	0	0	0	X	X	0	X	X
Ending delimiter	1	1	1	X	X	1	X	X

Figure 12.13 (the Priority fields are discussed below; Monitor Count is simply the Monitor Passed bit).

The second byte of the Control field is transmitted only with busy token frames. The Frame Format field indicates whether the data field contains ring-signaling information or data link information. If it indicates ring-signaling information, *all* stations are expected to read the Control field (shown in Figure 12.14).

The destination and source addresses are stored in the DST and SRC fields, respectively. The Data field is variable length, consisting of 0 or more bytes. The FCS is a 4-byte CRC and covers the entire frame between starting and ending delimiters. The free token frame consists of the two delimiters and the first byte of the CTRL field.

In the generic token ring description, all stations have equal opportunity

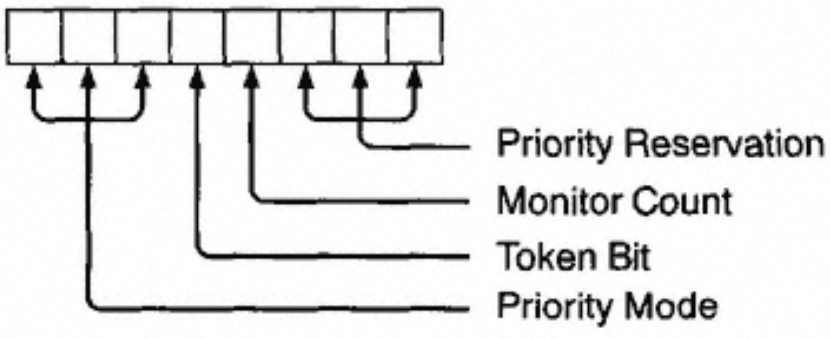

Figure 12.13
The format of the first CTRL byte.

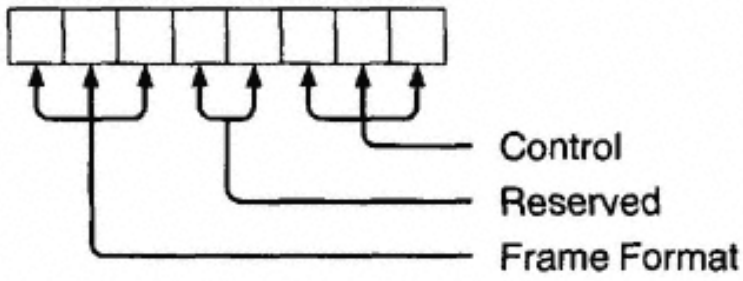

Figure 12.14
The format of the second CTRL byte.

to transmit any frame to any destination. This can mean that priority frames (such as voice or video) are competing with data frames that potentially could be delayed. To overcome this limitation, busy token frames are divided into two modes: those that can be transmitted asynchronously (periodically) and those that need to be transmitted synchronously (on a regular basis). To handle synchronous communications, each busy frame can be associated with one of eight levels of priority, as indicated by the priority bits in the CTRL field (0 is the lowest priority, indicating asynchronous mode).

Normally, the ring functions in asynchronous mode, with all stations having access to it. Periodically, when a high-priority station requires network access, the requested priority can be written into the Priority Reservation bits of a passing busy token. The station freeing the token is expected to copy these bits into the Priority Mode bits of the new free token. Stations with messages to be transmitted at the level of priority indicated in the Priority Mode bits may now do so, ensuring that higher-priority information can be sent. The priority does eventually come down, since the station requesting the higher priority is expected to return the token's priority to its original value after the synchronous communications have taken place. In the worst case, stations with asynchronous data to be sent can wait forever if all traffic is devoted to synchronous communications.

The frame status byte, FS, follows the final delimiter and is used by the destination station to convey status information back to the transmitting station through two status indication bits, A (acknowledgment) and C (framecopied) (R denotes reserved). Table 12.3 shows the format of the frame status byte. The interpretation of the A and C bits is given in Table 12.4.

The token ring operates at 4 Mbps.

The IEEE 802 committee has a token ring network standard: IEEE 802.5, which is compatible with the IBM Token Ring design. The addressing scheme follows that of the IEEE 802.3, discussed in Chapter 11.

Table 12.3 Format of the frame status byte.

R	R	C	A	R	R	C	A

Table 12.4 Interpretation of the frame status bits.

A	C	Meaning
0	0	Destination not responding.
1	0	Destination present but frame not copied.
1	1	Destination present and frame copied.

FDDI (Fiber Distributed Data Interface) is a 100-megabit-per-second (Mbps) fiber optic token ring standard developed by the American National Standard Institute (ANSI) committee X3T9.5, based upon the IEEE 802.5 token ring standard. An FDDI ring falls into the category of Metropolitan Area Network (MAN), since it can span distances of up to 200 kilometers. As a result of its high speed, it can also be used as a backbone network, interconnecting smaller, low-speed, local area networks (such as Ethernets or 802.5 token rings).

Physically, the FDDI network consists of two rings, one transmitting clockwise, the other counter-clockwise. Transmissions can occur on both rings (although the standard recommends that one ring remain in reserve), which gives an FDDI network an effective rate of 200 Mbps. Two classes of station exist: an A type, which connects to both rings, and a B type, which connects to a single ring. Up to 1000 stations can be joined to an FDDI network with a maximum distance of 2 kilometers between class A stations and 500 meters between class B stations. The relationship between class A and class B stations is shown in Figure 12.15.

As in 802.5, stations wait for the free token before transmitting. When a station has the token, it has a finite amount of time in which to transmit one or more packets. FDDI defines two packet types: synchronous, those a station is guaranteed to transmit; and asynchronous, those a station can transmit if there is time (there are eight levels of priority within asynchronous mode). Once a station has sent its last packet, it reissues the free token (FDDI operates as a multiple-token ring). It is responsible for removing its packets from the ring even though it no longer has the token.

One of the arguments for a bidirectional ring is reliability. Should both rings be cut for some reason, communications can continue by looping one ring back onto the other in the class A stations nearest the break, as shown in Figure 12.16.

There are presently two FDDI standards: FDDI-1, intended for data traffic; and a new standard, FDDI-2, intended to offer better support for both data and voice.

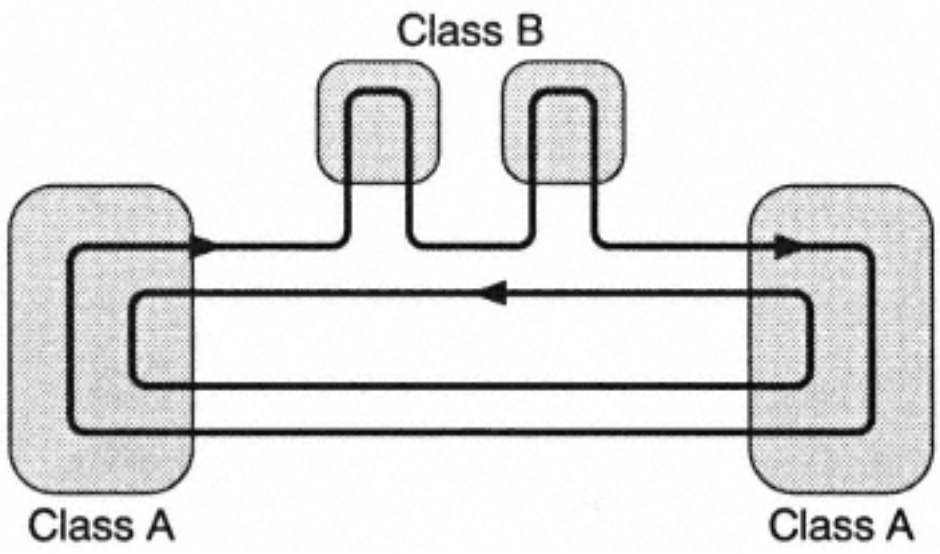

Figure 12.15
The relationship between FDDI Class A and B stations.

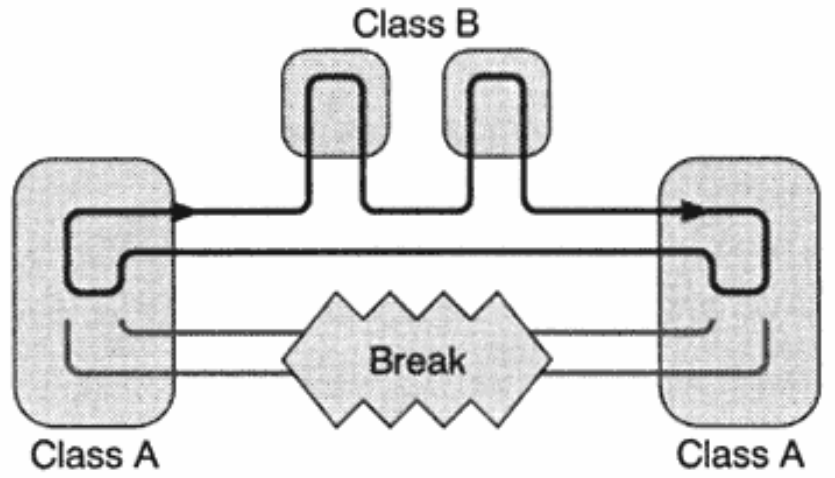

Figure 12.16
The bidirectional ring permits additional reliability.

12.3.2
The Slotted Ring

A slotted ring is similar to the token ring, except that there is a single circulating frame structure that contains space for data (anywhere from 2 to 32 bytes, depending upon the implementation). The frame has two states: free (indicating that a station can send data in this frame) or busy (meaning that this frame has data and should be read). As with the token ring, when the transmitting station receives its busy token frame, it frees the frame to allow another station access to the network.

Slotted rings typically have several circulating frames to improve throughput. The number of frames on the ring is controlled by the monitor and must be known by all stations to ensure that a transmitting station frees the correct frame.

Example

Probably the best known slotted ring to achieve any degree of commercial success was the Cambridge Ring, developed in the late 1970s at the University of Cambridge. The format of a Cambridge Ring frame (more commonly known as a minipacket) is shown in Figure 12.17 (all sizes are in *bits*).

The Start bit, with a value of 1, precedes all other bits in the minipacket and is used by the ring hardware as a synchronization bit to signal the start of the minipacket. The status of the packet then follows and indicates whether the packet is in use (Full, value 1) or available for use (Empty, value 0). The Monitor bit is set by the monitor to handle the circulating busy situation; this happens when a Full token passes the monitor and is cleared by the transmitting station.

Start (1)	Full/Empty (1)	Monitor (1)	DST (8)	SRC (8)	User Data (16)	Type (2)	Response (2)	Parity (1)

Figure 12.17
Cambridge Ring minipacket structure.

The destination, DST, and source, SRC, addresses are both 8 bits long. At most, 254 stations are allowed on a single ring; addresses `0x00` and `0xFF` are reserved. Two bytes of data are sent in the User Data field, while the Type field indicates the type of data.

The Response bits are sent by the transmitting station with the value 11; they are to be changed by the destination station and are interpreted as follows:

`11`the destination has not changed the response bits, which means that it does not exist.
`01`the destination has accepted the minipacket.
`10`the destination has accepted minipackets from another source.
`00`the destination is busy and cannot accept the minipacket at this moment.

The monitor station is a separate device; if it fails, the entire ring will shut down.

12.3.3
The Register Insertion Ring

The register insertion ring doesn't really correspond to any of the ring algorithms that have been discussed so far; instead, it functions more as a ring of store-and-forward stations.

Each station in a register insertion ring has two registers, each used to hold a variable-length frame: one from the upstream station and the other from the local station. Frames arrive at a local station as a stream of bits, which are copied into the upstream register. When the destination address bits arrive, the local station compares the address in the register with its own address. If the frame is addressed to the local station, it is copied to the station; otherwise, it is forwarded to the downstream station, one bit at a time (the upstream register acts as a queue, with the oldest bits being sent first and the most recent arrivals stored at the end).

Before the local station can transmit, the frame must be stored in the station's register. When the upstream register is empty, the station's ring hardware checks the station's register; if there is a frame to be transmitted, the bits in it are shifted onto the ring to the downstream station. Should bits arrive from the upstream station, the local station copies them into the upstream register while completing the transmission from the station's register.

12.4
The Commkit Token Ring

Although IBM manufactures a token ring, few PCs built by IBM (or any other company, for that matter) are supplied with token ring hardware as standard equipment. Nevertheless, as with the other networks studied in this

book, it is possible to write emulation software that illustrates the principles of ring networks.

Commkit is supplied with a token ring emulator, implemented with the following software:

`ringlow.c` the low-level token ring emulation software.

`ringdefs.h` the header file required by `ringlow.c`, containing the various data structures used by the ring network software.

`ringex.c` a high-level interface allowing experimentation with the token ring.

12.4.1
Hardware Requirements

The Commkit ring network software is written for PCs supporting two serial ports. A typical configuration consists of four PCs, each having two serial ports, with null-modem cables connecting the ports as shown in Figure 12.18. Note that unlike the bus network, serial port 1 *must* connect to serial port 2 since the ring algorithm stipulates that bytes arrive on port 2 and be forwarded on port 1.

12.4.2
Design

The ring network supplied with Commkit is a true token ring network in that before a message can be transmitted, the station with the message must first wait for the free token to appear. When the free token arrives, the station sets it to busy and sends the bytes in the message, one at a time, out serial port 1. All stations receive copies of the bytes making up the message as they circulate around the ring. Each byte arrives on serial port 2, a copy of the byte is taken, and the byte is then forwarded out serial port 1. The bytes are removed from the ring by the transmitting station, at which point a new free token is issued and allowed to pass to the next station in the ring. For example, if station B transmits a 4-byte message, the flow of information will be as shown in Figure 12.19 (the flow is counter-clockwise from station B).

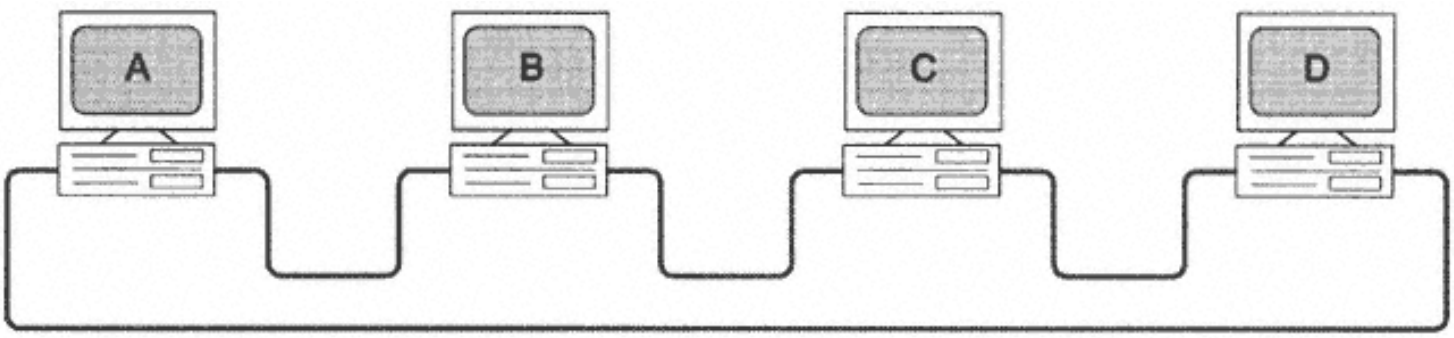

Figure 12.18
A Commkit ring network with four PCs.

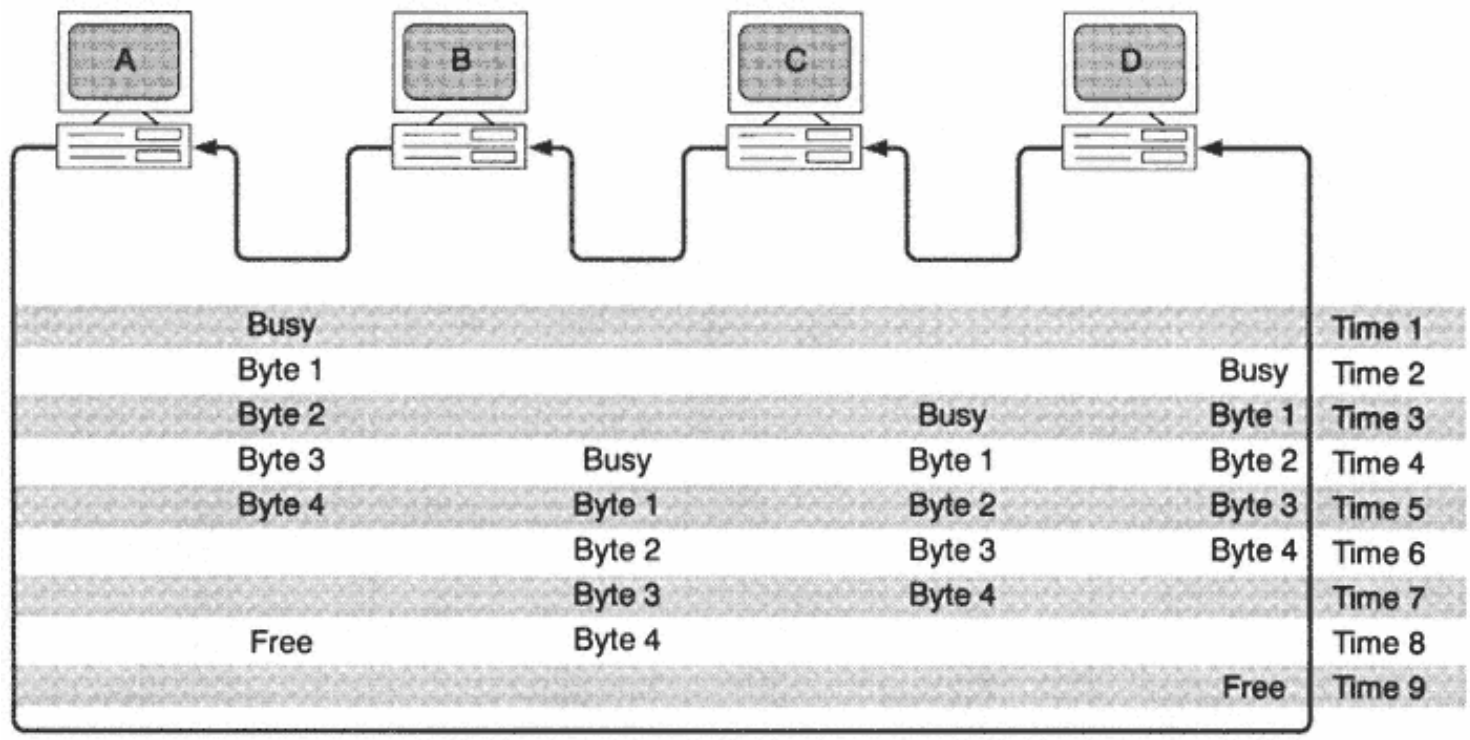

Figure 12.19
Information flow in the Commkit ring.

In the figure, at Time 1 station B puts a Busy token on the ring, which is received by station A. At Time 2 the Busy token has reached station D (after being forwarded by station A), and the first byte (Byte 1) arrives at station A. The Busy token is removed from the network at Time 4, but station B waits until the last byte (Byte 4) is received at Time 8 before putting the Free token on the network.

As with most other token rings, the Commkit ring network requires that one station be the monitor to ensure that the token is not lost or damaged. The monitor checks for the missing token and the circulating busy situations.

The Frame Structure

All information sent on the ring (whether a free token or a busy token and data) is sent in a variable-length frame. The general format of a Commkit ring frame is shown in Figure 12.20.

There are two valid frame structures:

Free token. A 3-byte frame that circulates past each station; it indicates that the network can be accessed. A station that has a message to transmit can seize the token, set it to busy, and transmit the message. Once the message has been sent, it is the responsibility of the transmitting station

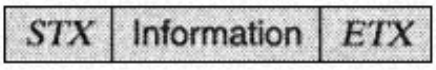

Figure 12.20
General format of a Commkit ring frame.

Figure 12.21
The free frame structure.

to place a new free token on the ring. The format of the free token frame is shown in Figure 12.21 (each field is one byte long).

STX and *ETX* delimit the token packet, while CTRL is the control byte. The free token CTRL byte has a value of `0x04`.

Busy token. A variable-length frame that can hold up to 64 bytes of transparent data. The framing characters are *STX* and *ETX*. The overall frame structure is shown in Figure 12.22 (all sizes shown in parenthesis are expressed in bytes).

The control byte, CTRL, contains one of two values:

`BUSY`. The frame contains data for a specific station (as indicated by the DST field). A control field set to `BUSY` has a value of `0x08`.

`MNTR_PASSED`. The frame is busy (`0x08`) and has passed the monitor (`0x10`).

DST and SRC are the destination and source addresses of the frame, respectively; both are one byte long.

To avoid confusing data *ETX* with the end-of-frame delimiter *ETX*. the ring network software prefixes all *ETX* and *DLE* bytes with a *DLE*. When bytes are removed from the network for local storage by a station, the prefix *DLE* character is discarded; however, the frame contents are not modified as they are forwarded by a station.

The CRC is generated by the transmitting station using the CRC-16 algorithm and covers the entire frame except for the frame delimiters and the control field. The control field is not included because its value changes as it passes the monitor station.

The Emulator States

A ring station is always in one of three states: idle, transmitting, and receiving. A station is considered to be in the idle state if it has nothing to transmit although a free token has been detected. If the token is busy, the station enters the receive state, copying and forwarding each byte until the end-of-frame byte is detected. The station enters the transmit state if the token is free and a message is waiting to be sent. The token is first set to busy and

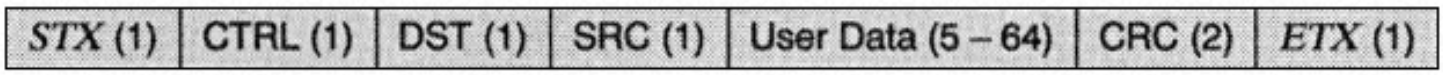

Figure 12.22
The busy frame structure.

the message is then transmitted, a byte at a time. Upon completion of the transmission, a free token is put back on the network.

The Idle State. A station is in the idle state if it has just forwarded a free token (either because it has no data to send or because it has just finished transmitting). All stations enter the idle state as soon as they are initialized. What a station does while idle depends upon whether or not it is the monitor.

If the station is the monitor, it is responsible for ensuring that there is always a token on the ring. This is achieved by running a timer in the background which periodically causes the monitor to check whether token frames (free or busy) are still circulating. If none have been detected, a free token is put on the ring. The ring monitor state machine is illustrated in Figure 12.23.

The monitor initially waits for an *STX*; when it receives one, the byte is forwarded and the monitor waits for the control byte. Anything other than

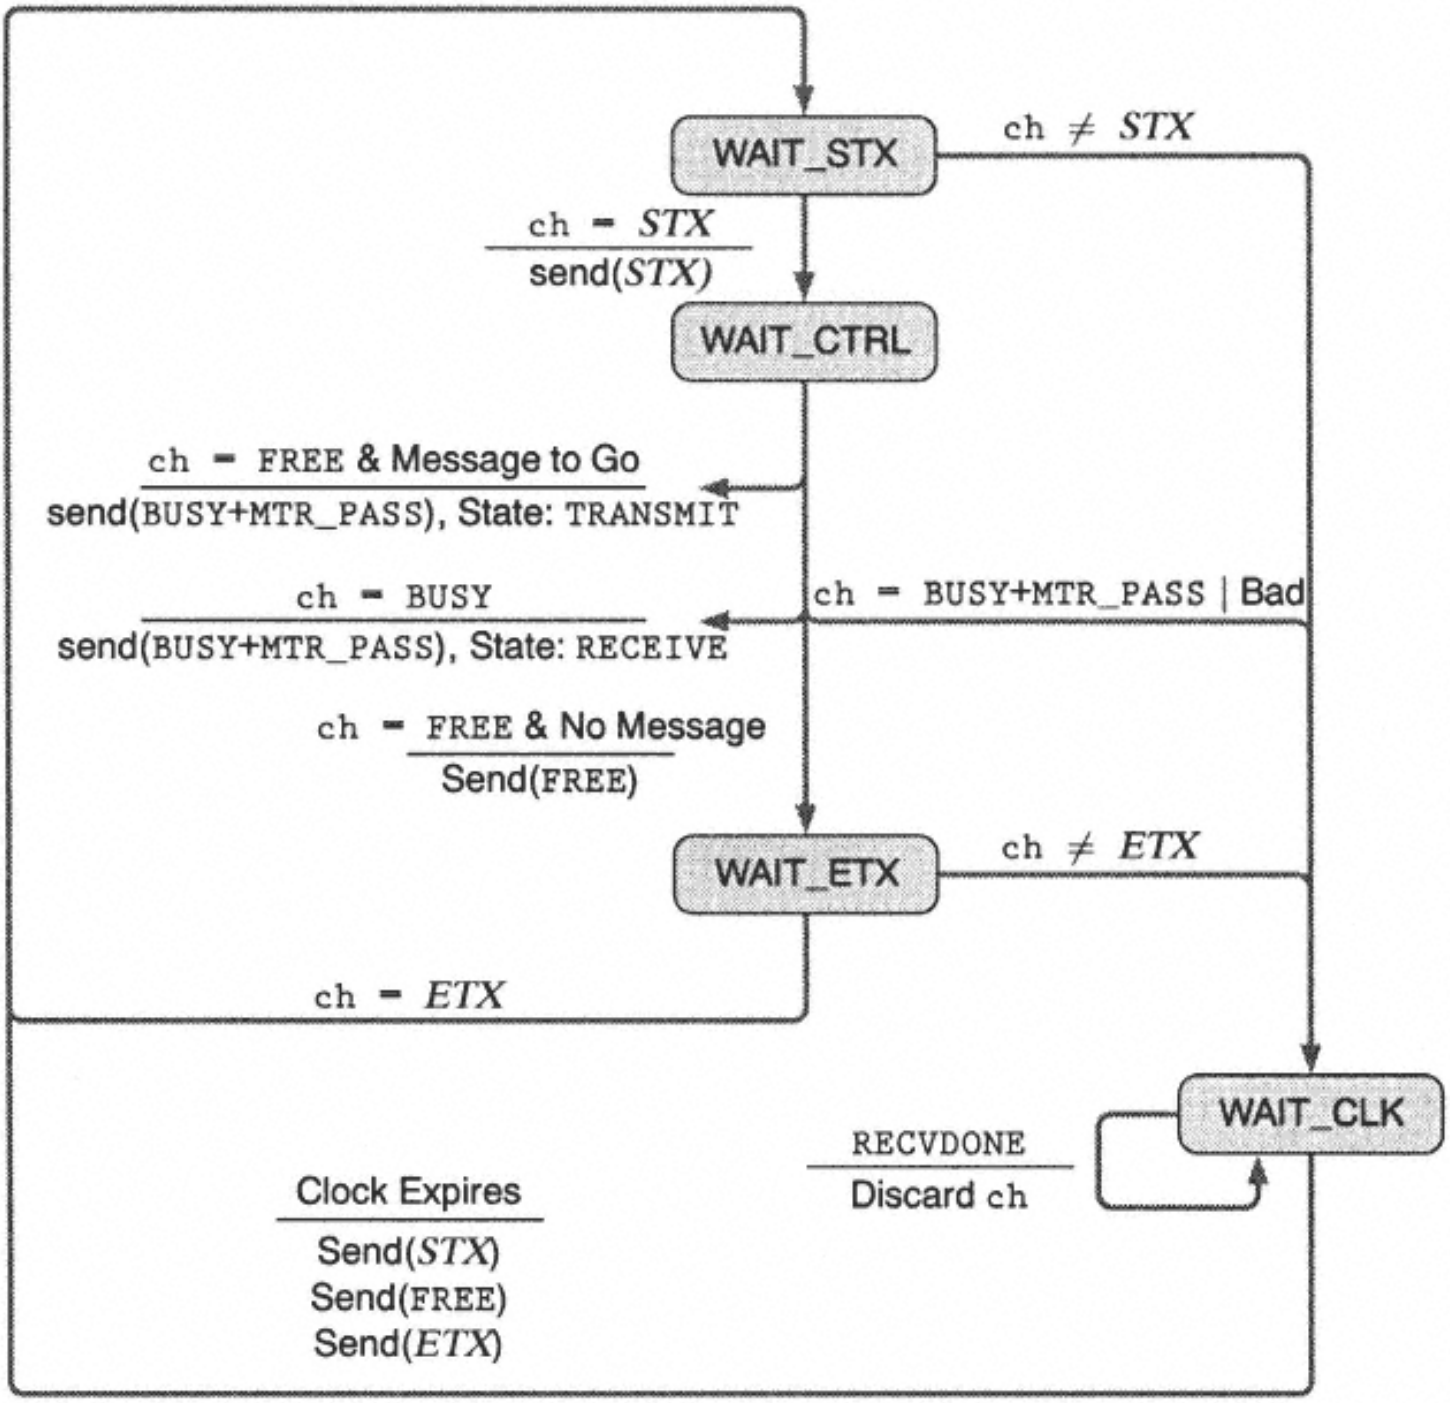

Figure 12.23
Ring monitor state diagram.

an *STX* causes the monitor to enter the error state, WAIT_CLK. When the control byte is received, it is examined by the monitor:

If the byte indicates a FREE token and the monitor has a message to transmit, the control byte is set to BUSY and MTR_PASS (signaling that the monitor has detected a busy token). The monitor enters the transmit state to allow the transmission of the addresses, the message, and the CRC.

If the token is BUSY, the monitor changes the control byte to include MTR_PASS as well as BUSY. The control byte is then forwarded, and the monitor enters the receive state to wait for the remainder of the busy token frame.

If the monitor has no messages to transmit and the control byte indicates a FREE token, the control byte is forwarded and the monitor waits for an *ETX*.

If the next byte is an *ETX*. the monitor forwards it and waits for an *STX*. Anything other than an *ETX* causes the monitor to enter the error state, WAIT_CLK.

If the control byte indicates a BUSY token as well as the token having passed the monitor (MTR_PASS), or if the control byte is unrecognizable, the monitor enters the error state WAIT_CLK.

All monitor idle states are associated with a timer. Should a byte fail to arrive at the monitor within a prescribed period of time, the monitor assumes the worst and reissues a new free token on the ring. If a bad or unknown byte is received by the monitor, the monitor enters the WAIT_CLK state; all subsequent bytes are removed from the ring, and the monitor issues a new free token.

When the monitor is first initialized, a free token frame is sent and the timer is started. If the timer expires before anything is received, another free token frame is issued.

Although all stations support the monitor software, only one can be the monitor at any moment. If a station is not the monitor, it uses the state machine shown in Figure 12.24 for the idle state.

Except for a number of housekeeping states devised to ensure that the token is circulating, the two state machines are identical. For example, if the monitor has a message to transmit, it must wait for a free token exactly as a nonmonitor station does. The additional monitor states enforce the free token frame structure of *STX*-CTRL-*ETX*. The nonmonitor station leaves the correction of an invalid frame to the monitor.

The Transmit State. The transmit state consists of *two* state machines. The first is for the transmission of data on the ring, and the second is for the removal of the data from the ring once it has completed its journey around it. The transmission state diagram is shown in Figure 12.25.

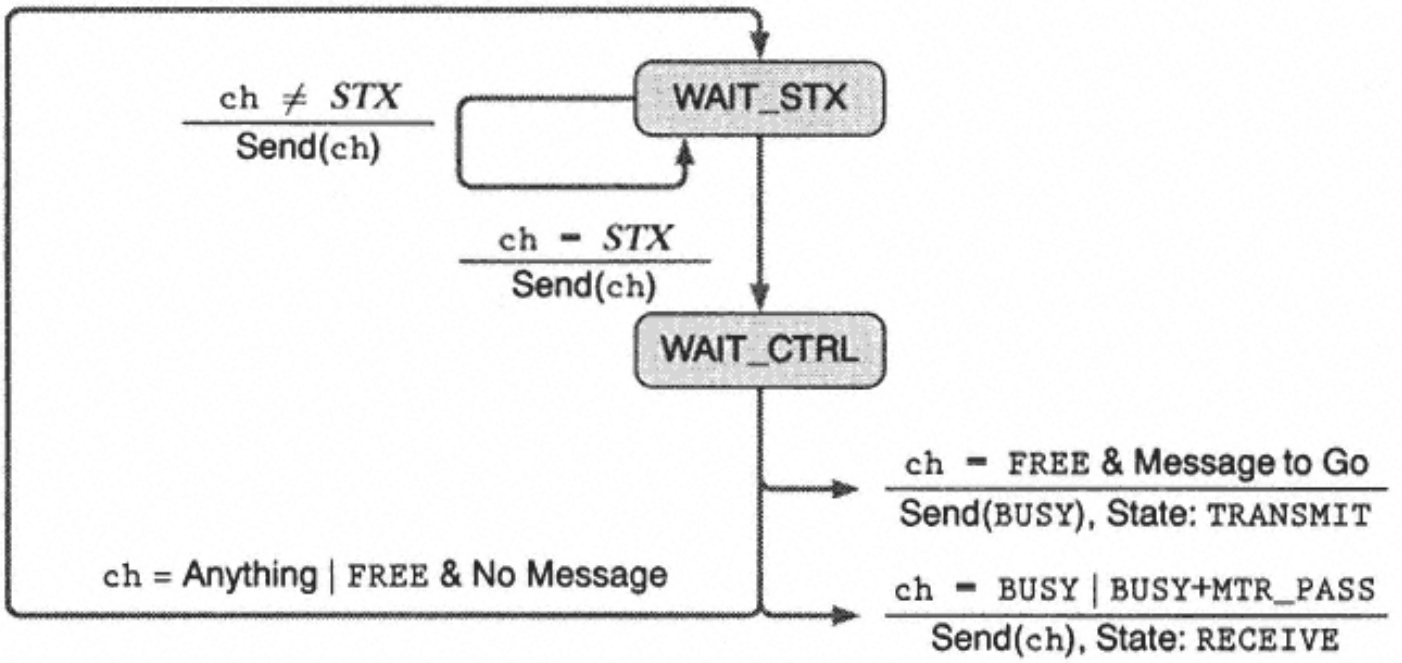

Figure 12.24
Nonmonitor idle state.

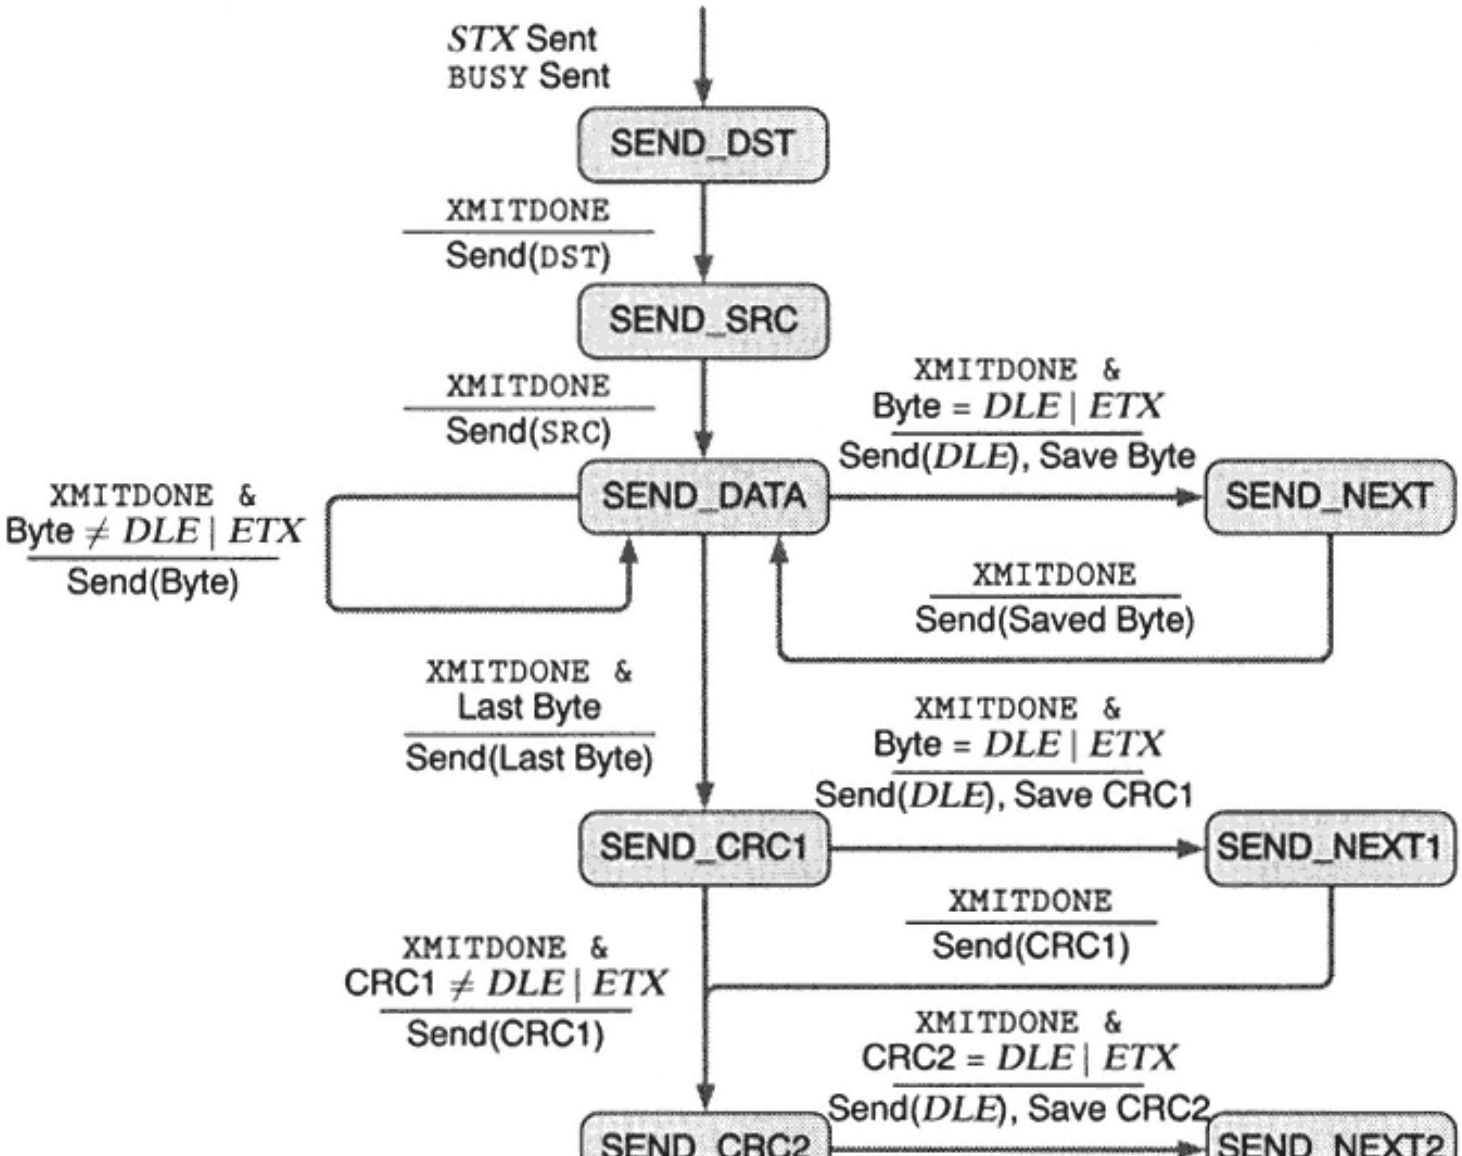

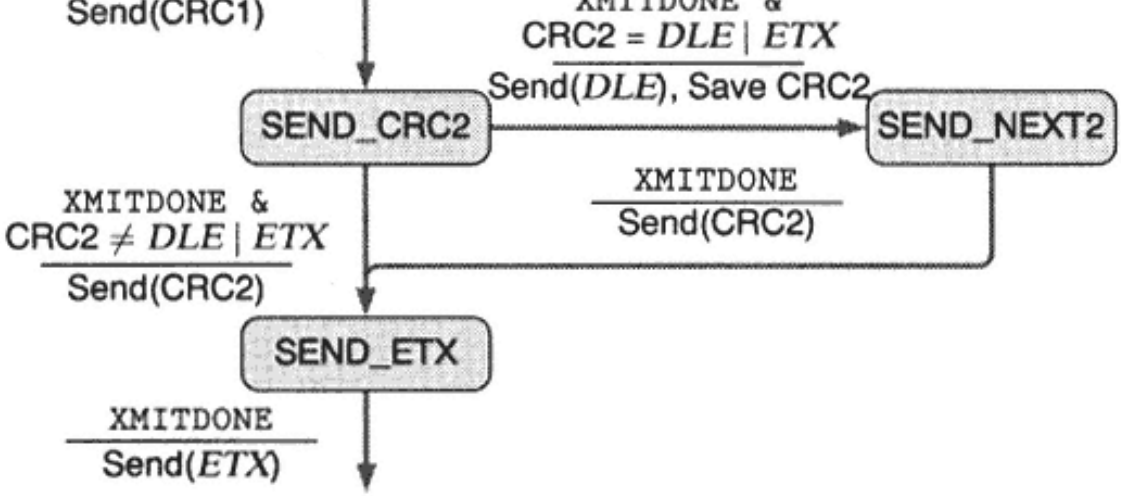

Figure 12.25
State diagram for transmitting a frame.

The *STX* and the control field (set to BUSY) are already sent by the time the transmit state machine is initiated. The transmit state machine is responsible for sending the destination and source addresses, the bytes in the message, the two CRC bytes, and the final *ETX*. Additionally, *DLE* insertion takes place whenever a *DLE* or an *ETX* is found in the data stream or as a CRC byte. The CRC includes all transmitted bytes *except* for the control field, any *DLE* prefix bytes, the CRC bytes, and the frame delimiters.

The transmitting station is also responsible for removing its packet from the network. A second state machine, shown in Figure 12.26, runs in parallel with the transmission state machine, removing the packet from the ring.

The removal of a packet from the ring involves taking the remainder of the previous free packet from the ring (an *ETX*) and then waiting for the incoming *STX*. Once the *STX* is detected (and removed), all subsequent bytes that make up the packet can be removed. Should an *ETX* be detected in place of the control field, the destination address, or the source address, it is assumed that the transmission has been aborted, probably by the monitor. *DLE* insertion must be recognized to avoid confusing a data *ETX* or a CRC *ETX* with the end-of-packet *ETX*. CRC calculation is not performed.

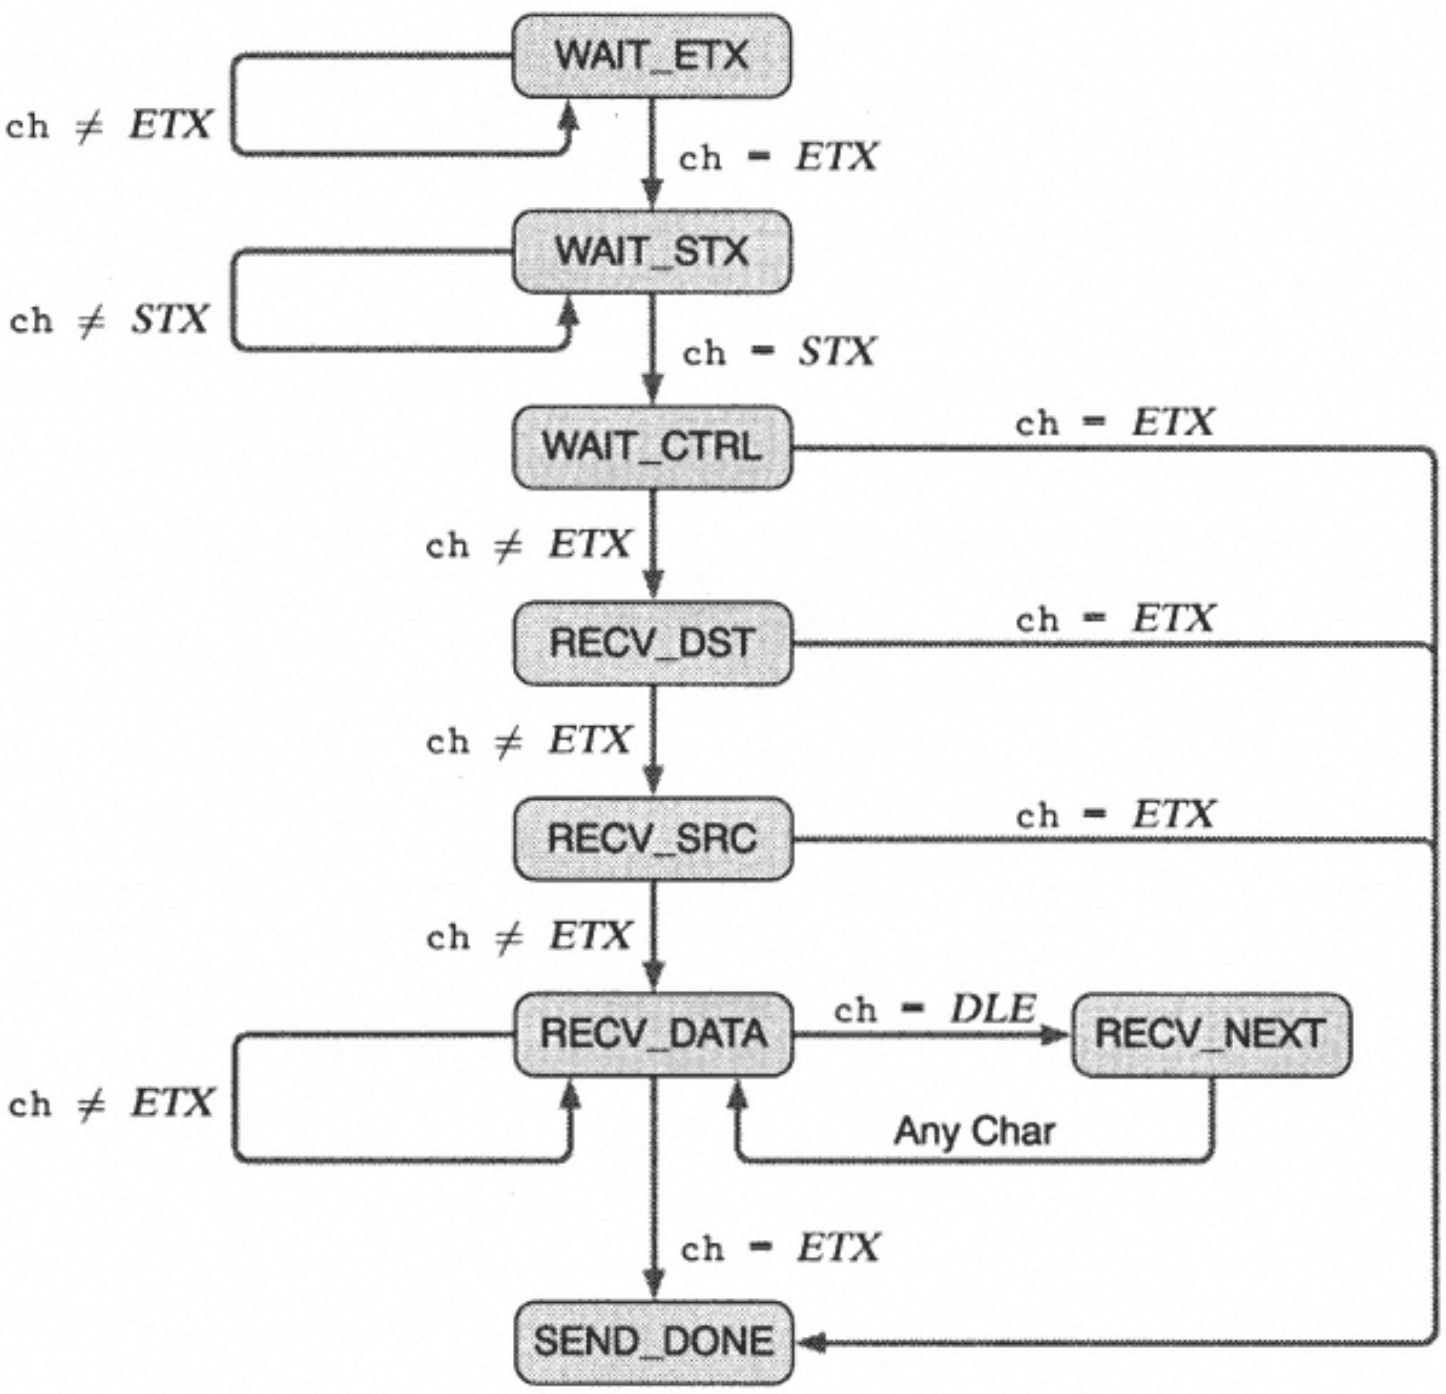

Figure 12.26
State machine for removing frames from the ring.

The Receive State. The receive state is entered after a station receives a packet with the token bit set to BUSY in the control field. It is responsible for copying the message from the network and storing it. Each byte received must be forwarded to allow all the stations on the network access to the packet. The state machine for a station receiving a packet is shown in Figure 12.27.

Upon detection of a BUSY token, the station waits for the remainder of the frame. Data is removed (and forwarded) in state WAIT_DATA, while *DLE* removal is handled in state WAIT_NEXT. (Note that *DLE* removal only affects the data to be stored; the *DLE* character is still forwarded.) When the end-of-frame delimiter *ETX* is found, both the CRC and the destination address are examined; if the CRC is 0 and the destination address is that of this station, the message is kept; otherwise, it is discarded.

12.4.3
The Token Ring Implementation

The Commkit token ring is an implementation of the state machines described in the previous section. The emulator, found in `ringlow.c`, is interrupt driven; serial port and clock interrupts are passed through `low_level()` to `ring_emulator()`. Keyboard interrupts are forwarded to the foreground process (queue APPLICATION), while clock interrupts are forwarded to the background process (queue BACKGROUND_1) HZ times a second. Messages to be sent on the network are supplied by any process to the ring via the queue RINGNET (a redefinition of SP2IH). MSG_AVAIL is ignored, since the ring software checks the RINGNET queue only when a free token is received.

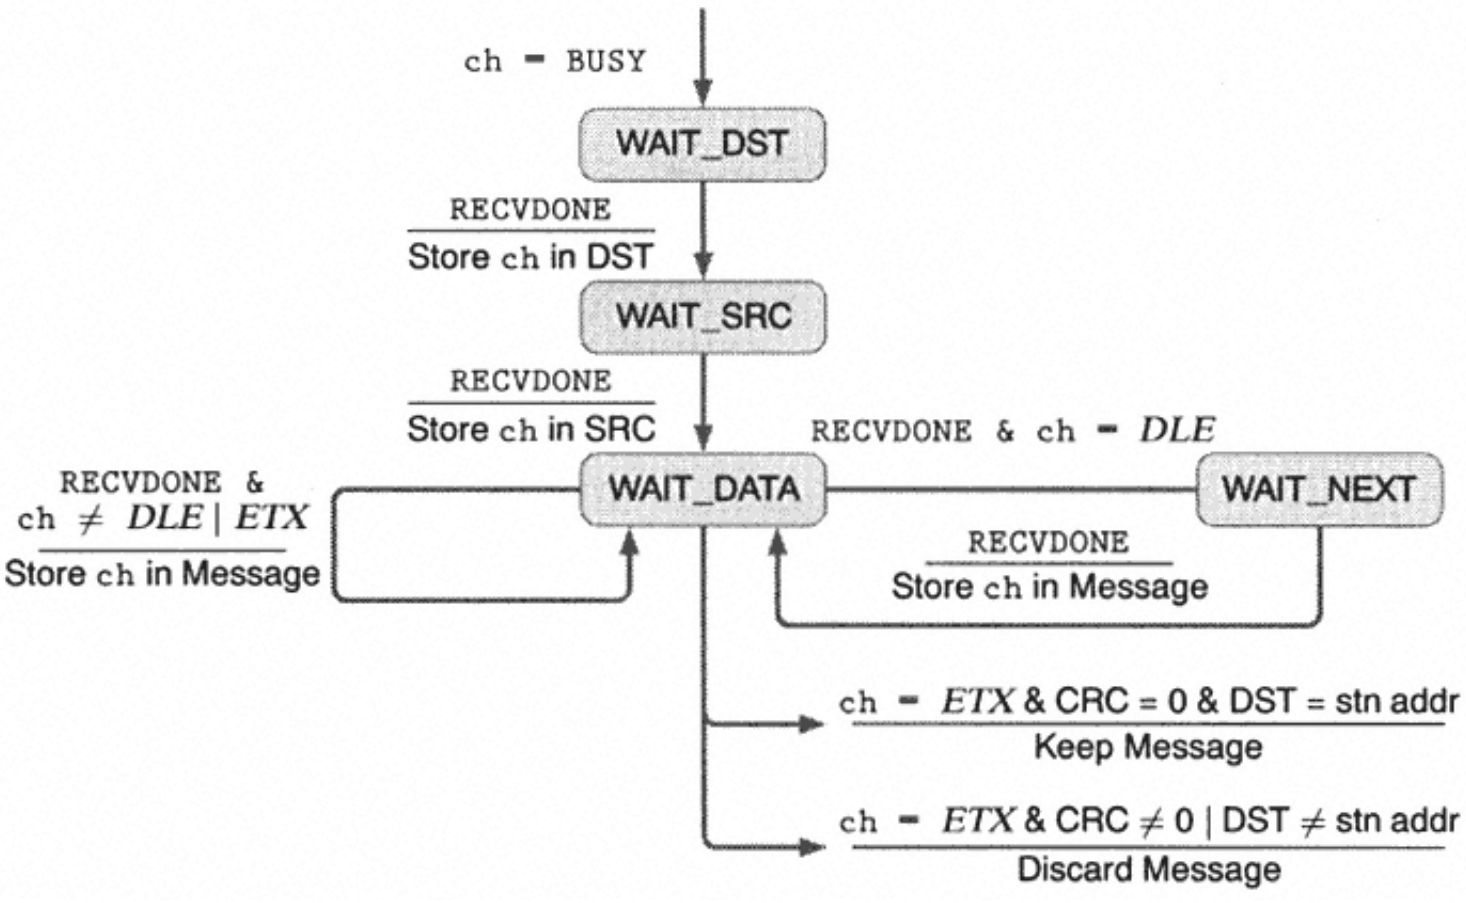

Figure 12.27
State diagram for receiving a frame.

`Ring_emulator()` is called after a serial port interrupt (either `XMITDONE` or `RECVDONE`) occurs. Each byte received may (or may not) be transmitted to the next station, depending upon the state of the emulator. `Ring.char_to_go` (found in `ringdefs.h`) indicates whether a character can be transmitted. The fact that the byte just received can be transmitted does not necessarily mean that it will be transmitted immediately: other bytes may be waiting to be transmitted, or a byte may be in the process of being transmitted (indicated by `ring.xmit_active`). Before attempting to transmit the byte, `ring_emulator()` checks the queue `SP1IH`; if the queue is not empty, the byte must be put on `SP1IH` (the `SP1IH` queue acts as an elastic buffer, holding characters waiting to be transmitted).

In the various emulator states, whenever a byte is to be transmitted, it is stored in `ring.to_go` and `ring.char_to_go` is set `TRUE`. The three ring states (indicated by `ring.state`) and their related procedures are as follows:

The Idle State. Entered when `ring.state` is equal to `IDLE.Ring_idle()` implements both the monitor and monitorless ring idle state machines in a single procedure. Whenever a free token is received, the `RINGNET` queue is checked, and the presence of a message causes the token to be set `BUSY` (or both `BUSY` and `MONITOR`, if the station is the monitor); the `ring.state` changes to `TRANSMIT`. If a `BUSY` or `MNTR_PASSED` token is received, it is flagged to be forwarded and the station enters the `RECEIVE` state. Should the monitor detect anything amiss in any state, it enters the `WAIT_TIMER` state, removing all data from the ring until the timer, `ring.token_timer`, expires. `Token_timer` is decremented in `check_for_token()`, which is called once a second after a clock interrupt in `low_level()`.

The individual state within the idle state is indicated by `ring.r_state`.

TheTransmission State (indicated by `ring.state` having a value of `TRANSMIT`). Entered if a free token has been received and a message is to be sent. It is implemented as two procedures: `send_byte()` (called after an `XMITDONE` interrupt) and `clear_line()` (called after a `RECVDONE` interrupt).

`Send_byte()` transmits the various characters making up the packet: the destination address, the source addresses, and the bytes in the message (the overall message structure is defined in `ring.x_msg`). `Send_byte()` implements the transmission state machine described in the previous section; the specific transmission state within `send_byte()` is maintained in `ring.x_state`. If a byte requires a *DLE* prefix, `send_byte()` calls `send_dle()`, which "remembers" the byte to be sent as well as the next state.

Bytes are removed from the ring by calls to `clear_line()` after each `RECVDONE` interrupt. `Clearline()` implements the byte-removal state machine described in the previous section (the current state is indicated

by ring. r_state); the incoming bytes are not checked, nor is the value of the CRC. When ring. x_state is equal to ring.r_state (i.e., they both have the value SEND_DONE), a free token is put on the ring (by a call to send_token()).

The Reception State. Whenever a BUSY or MNTR_PASSED token is received, ring.state changes to RECEIVE. The receive state machine described in the previous section is implemented in recv_byte().

Recv_byte() accepts bytes from the ring whenever a RECVDONE interrupt is signaled. The first two bytes are taken to be the destination and source addresses, respectively. The remaining bytes are read from the ring, stored in ring. r_msg, and marked for transmission. The specific receiving state is indicated by ring.r_state. Upon detection of the end-of-frame *ETX*. the destination address and the CRC are inspected: if the local station is the intended destination, the message is sent to queue BACKGROUND_1; control then returns to the idle state. If the global variable promiscuous is TRUE, all valid frames are accepted and sent to queue BACKGROUND_1, regardless of destination.

Errors such as the packet being too long or a bad CRC cause the reception to be aborted, and ring.state is changed to IDLE.

All of the routines described are common to both monitor and nonmonitor stations; they are distinguished by the value of bus.monitor (TRUE if the station is the monitor, FALSE otherwise).

12.4.4
High-Level Software

The token ring software performs two basic operations: taking messages off the RINGNET queue, turning them into packets, and transmitting them when the opportunity arises; and copying messages from the network and forwarding them to the BACKGROUND_1 queue. To facilitate ring network experimentation, ringex.c contains a number of processes through which the user can communicate with the ring network:

do_display(). Responsible for displaying messages sent by the token ring emulator to queue BACKGROUND_1. When a message from RINGNET is available, do_display() takes the message from the queue and displays it on line 10 enclosed in angle brackets. Messages from the clock interrupt handler are read but ignored.

do_ring_ex(). The foreground process, accepting characters sent by the keyboard interrupt handler, displaying and storing them in array buffer; it is assumed that the data is entered as a valid ring packet (i.e., the destination address, the source address, and up to 64 bytes of data). Upon receipt of a carriage return, buffer is forwarded to the token ring queue RINGNET for eventual transmission. A CTRL-C causes control to return to MS-DOS (the ring ceases to function at this point).

12.5 Using the Commkit Token Ring

12.5.1 Compiling the Token Ring

The token ring source code `ringex.c`, `ringlow.c`, and `ringdefs.h` can be compiled and linked with `commkit.obj`, `srcalls.obj`, and `crc.obj` to make `ringex.exe`:

```
C:\> make ringex.exe
```

If changes are made to `ringex.c`, the `make` utility recompiles only `ringex.c`, relinking to the existing object modules.

The Commkit diskette is supplied with an executable version of `ringex.exe`.

12.5.2 Running the Token Ring

To run a compiled version of the token ring, type `ringex`, followed by the line speed (50, 300, 1200, 2400, 4800, or 9600 bps) and the station identifier (any character) after the DOS prompt. For example, for the network to run at 1200 bps and this station to have address `E`, type the following:

```
C:\> ringex 1200 E
```

The ring network differs from the bus network in that one station must be the monitor; the monitor station is identified by `M` after the station identifier. The token ring will not function if there is no monitor station. For example, to make station `S` the monitor station, type

```
C:\> ringex 1200 S M
```

If the station is not the monitor station, simply ignore the `M` or type any other character. As with the other Commkit modules, if the line speed is not recognized or the format of the command line is not adhered to, an error message is generated.

Each PC on the ring must have two serial ports. The PCs are connected by null-modem cables linking serial port 1 with serial port 2. The ring will not function if the ports are not connected in this manner.

Once `ringex.exe` is running, a message can be passed between stations by typing the destination station's address, the source station's address, and then the message. For example, to inquire what the weather is like over by station `S`, one could type:

```
SEWhat's the weather like over there?
```

At station E, the message would appear part way down the screen enclosed in angle brackets:

```
<SEWhat's the weather like over there?>
```

Any station, including the monitor, can transmit and receive messages. To get a better understanding of how a ring network functions, place a line analyzer between two of the stations.

Control can be returned to MS-DOS by typing CTRL-C or CTRL-ALT-DEL. Unlike the other networks examined in the book, once a station is removed from the ring, all communications cease.

12.6 Summary

Local area networks allow device sharing among all stations. Networks such as the CSMA/CD bus can potentially restrict network access (for example, because of exceptionally heavy network traffic), since access is nondeterministic. By changing the topology (either logically or physically) to a ring, it is possible to develop a local area network that permits deterministic access. In other words, there is an upper boundary on the length of time a station must wait before accessing the network.

The solution adopted by both the token bus and the token ring is to permit network access only when the station is in possession of the "free" token. When a station receives the free token (from its predecessor, either logically or physically upstream from the station) and has a message to be sent, a transmission can take place. Once the station has finished transmitting, a new free token is placed onto the network, permitting network access to the downstream stations.

The token bus differs from the token ring in a number of areas; the most obvious is the topology: one is a bus and the other is a ring. Another difference involves network error recovery methods: in the token bus, recovery from network errors is distributed among all machines, while in the token ring the monitor station maintains control over the ring.

The token ring is by no means the only physical ring structure possible; a variation is the slotted ring, in which the network supports a single frame structure consisting of a control field with space for the token and space for the bytes in the message.

12.7 Exercises

1. (B) Set up the 3-station ring network shown in Figure 12.28; the serial port connections are shown in Table 12.5. Now, answer the following questions:

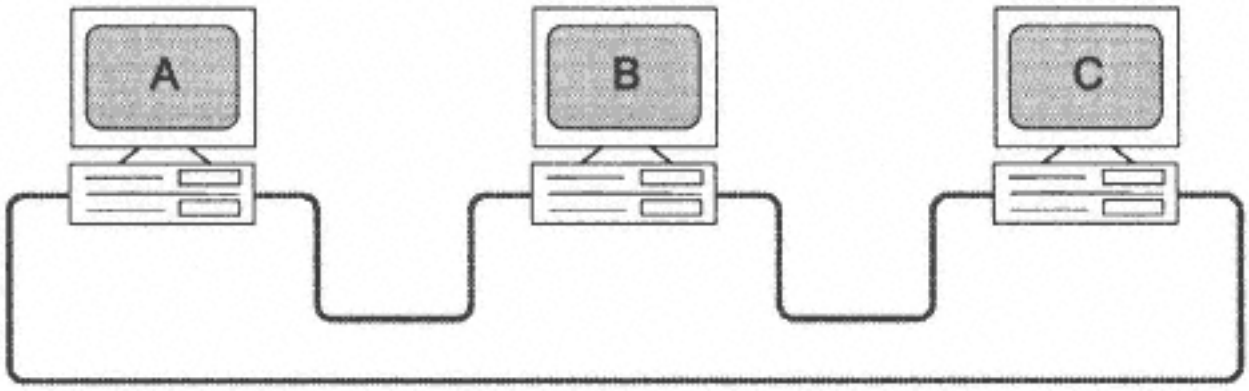

Figure 12.28
Station configuration for Exercise 1.

Table 12.5 Serial port connections for Exercise 1.

Station	Port 1 (to)	Port 2 (from)
A	C	B
B	A	C
C	B	A

a. If A is the monitor station, B is the analyzer, and C is inactive, what appears on the line analyzer? Why?

b. If C is the monitor station, B is the analyzer, and A is inactive, what appears on the line analyzer? Why?

c. If A is the monitor station, B is the analyzer, and C is an active ring station, what appears on the line analyzer? Send a message from A to C; explain what appears on the analyzer. Send a message from C to A; is there a difference in the value of the token? Explain.

d. If C is the monitor station, B is the analyzer, and A is an active ring station, what appears on the line analyzer? Send a message from A to C; explain what appears on the analyzer. Send a message from C to A; is there a difference in the value of the token? Explain.

e. If a station sends a message to a nonexistent station (say D), is the message transmitted and carried by the ring network? Explain.

Note that when testing the ring network with the line analyzer, the characters shown in Table 12.6 appear on the analyzer screen.

2. (B) Why do communications cease if a Commkit token ring station is removed from the network? Why shouldn't this happen on a commercial token ring?

3. (B) What happens if two (or more) stations are set up as the monitor? What happens if a message is sent by a third station? The line analyzer can help solve this problem.

4. (B) What happens if all stations are set up as nonmonitor stations? If communications cannot take place, explain why this is so. How can the problem be rectified?

Table 12.6 Line analyzer interpretation of Commkit frame characters.

Symbol	Meaning
	Beginning of frame.
	Free token indication.
↑	Busy token (Passed Monitor bit set).
	Busy token (Passed Monitor bit cleared).
	End of frame.

5. (B) Set up the line analyzer and a single token ring station (as the monitor station). Connect the ring network's serial port 1 to either of the line analyzer ports. What happens?

6. (B) Place the line analyzer between two stations on the token ring and watch the traffic. How does the ring's traffic differ from that on the bus network? Why is the ring said to be active while the bus is described as passive?

7. (C) The ring emulator does not handle two (or more) stations transmitting simultaneously: The packet is simply removed by each transmitting station. Modify the ring emulator software to handle multiple tokens.

8. (B) In the introduction to this chapter, the ring network was described as being more efficient than the bus network in some situations. Devise a set of experiments to compare the access times and overheads of the bus network and the ring network.

9. (C) Modify the Commkit token ring so that the duplicate token situation is handled; that is, two stations on the ring are attempting to transmit simultaneously, and the one with the lower address value stops its transmission and allows the other station to proceed.

10. (C) If the monitor station should cease to function, the remaining stations must bid among themselves to determine which station is to become the new monitor. Extend the ring software to allow the remaining stations to bid to become the monitor when the old monitor fails. Remember, the Commkit monitor station *cannot* be physically removed from the network, since it is still required to forward each byte on the ring. Therefore, write the software so the monitor station can be *logically* switched out of the ring (for example, reserve a key to signal the monitor software to shut down).

One straightforward technique is to set `ring.monitor` to `FALSE` when the user types CTRL-Z (ASCII *SUB*). Then break the ring by momentarily pulling a null-modem cable out of one of the serial ports. Ideally, the token will be either damaged or removed entirely from the ring and there is now no monitor.

11. (C) The Commkit ring network is a token ring. Modify the ring software so that it supports a slotted ring.

12. (C) If a channel on a ring network is broken, the circulation of the token stops. Robust ring networks (such as FDDI) allow tokens to circulate in both directions so that if a break is detected, the loop is still maintained. Modify the Commkit token ring so that a cable break can be overcome by circulating the token in the reverse direction when it arrives at the station where the break has occurred. The design of this is fairly straightforward, since the RS-232 cable connecting the stations allows full-duplex communications; a backward channel already exists in hardware.

13. (A) When the transmitting station removes the packet from the ring, no CRC calculations are performed. Are CRC calculations useful when the data is removed from the ring? What benefits, if any, does this give to the transmitting station?

PART 5
INTERNETWORK COMMUNICATIONS

Until now, all network communications have been intranetwork; that is, they take place within a single network. However, there are times when an application on a machine on one network must communicate with a machine on a separate network; consider the following examples:

A company consists of a number of separate departments, each with its own local area network. If electronic information, such as reports or personnel data, is to be exchanged between the various departments, there must be a mechanism to transfer it between the different networks.

A group of universities and colleges may all contribute to purchase a "supercomputer," which is sited on a local area network at one of the schools. Access to the supercomputer from research machines on local area networks at other institutions may require the interconnection of two or more of the local area networks.

The above are examples of applications that require internetwork communications, or communications *between* networks, which take place across a number of interconnected networks. The various networks that make up the internetwork are referred to as an internet or catenet (short for "concatenated network").

Not surprisingly, an internet spanning a region or even a country exhibits many similarities to a wide area network (WAN). However, internetwork communications are different from WAN communications for a number of reasons. First, a WAN is physically a single network, whereas an internet is made up of many distinct networks (that may include wide area networks). Second, the WAN supports a common set of data link protocols, shared by the nodes that make up the network; in an internet, different networks will

probably support different data link protocols. Third, a single addressing structure is used within a wide area network, ensuring that each node is uniquely identified; in an internet, each network supports its own addressing scheme, with no guarantee that the addresses are not duplicated on other networks.

Part 5 examines two issues common to the development and implementation of any internet:

How an application on one network identifies the intended destination on a remote network.

How the various networks involved in an internet are interconnected.

Connecting one network to another does not ensure that the two networks can communicate. If nothing else, an application on one network must be able to identify the intended destination service. Chapter 13 examines some of the issues surrounding the development of layered network architectures to support internetwork communications. Commkit's version of sockets is used to illustrate some of the problems associated with the development of an internetwork architecture.

If all networks making up an internet supported the same protocols and electrical characteristics, interconnecting networks would be a simple task. Since forcing a single network standard on the groups making up the internet may well be impossible (some groups may have already purchased their networks), a mechanism is required to handle protocol conversion when messages flow between the various networks. These mechanisms, known generically as gateways, are presented in Chapter 14 and illustrated by examining the interconnection of two Commkit networks.

13
Layered Architectures

13.1
Introduction

With few exceptions, most of the work done thus far on the various Commkit networks has

Required the user to explicitly build the network packet, identify the destination and source addresses, and, in some cases, enter network-specific control information. Although this is useful in teaching network functioning, it is unrealistic to expect the everyday user (other than, perhaps, someone working on a research project) to perform such tasks. Imagine the user's enthusiasm if, every time a message on an 802.5 network is to be sent, the destination's 48-bit address must be typed explicitly.

Assumed that the source entity is a process that accepts keyboard characters and forwards them to a remote DTE, where the destination process is responsible for displaying the characters. Many operating systems support multiple processes on a single host, meaning that a message arriving at a host for "the process" has little meaning unless the destination process can be explicitly identified.

This chapter considers how communication systems overcome the problems of identifying entities and hiding network-specific features using layered software, to produce what is commonly known as a layered architecture.

13.2
Background

13.2.1
Identifying Remote Entities

Operating systems that support multiple processes require mechanisms for identifying individual processes. For example, in Commkit each process is

associated with a unique identifier (such as `APPLICATION, BACKGROUND_1`, and so on) and a queue. A source process sends a message to a destination process by supplying the message and the identifier of the destination process to the queue management software using the `send()` primitive. Similarly, when a Commkit process is to receive a message, the `recv()` primitive is invoked, supplying the queue management software with the process's identifier. The queue associated with the identifier is then accessed, and the first available message on the queue is returned to the process.

This model can be extended to the transmission of messages between hosts. That is, a frame containing the process identifier of the destination process can be sent to the destination host. The destination address then consists of two parts: the destination host and the destination process. Upon receipt of the frame, the reception software on the destination host performs a `send()` to the destination process specified in the destination address. Overall, the communication software functions like a multiplexer: messages to various destinations are multiplexed into separate frames and transmitted on the network, while frames are taken from the network and demultiplexed into messages using the process identifier. A typical sequence is shown in Figure 13.1.

Using process identifiers to identify remote processes works successfully as long as

There is a one-to-one mapping between a process and its process identifier. In many operating systems supporting multiple processes, a pool of identifiers is recycled among the existing processes. Over a period of time, an identifier may be associated with a number of different processes; as processes terminate, the identifier is freed to be used by a newly-created process. If process identifiers are assigned dynamically, there is no guarantee the destination process will be associated with the "correct" one.

All identifiers are the same format. Different operating systems often use different identifier formats and sizes. This can lead to problems in a network of heterogenous machines. For example, if a 16-bit identifier size is chosen and an operating system that uses 32-bit identifiers is attached to the network, how are 32-bit identifiers to be represented? Similarly, if a

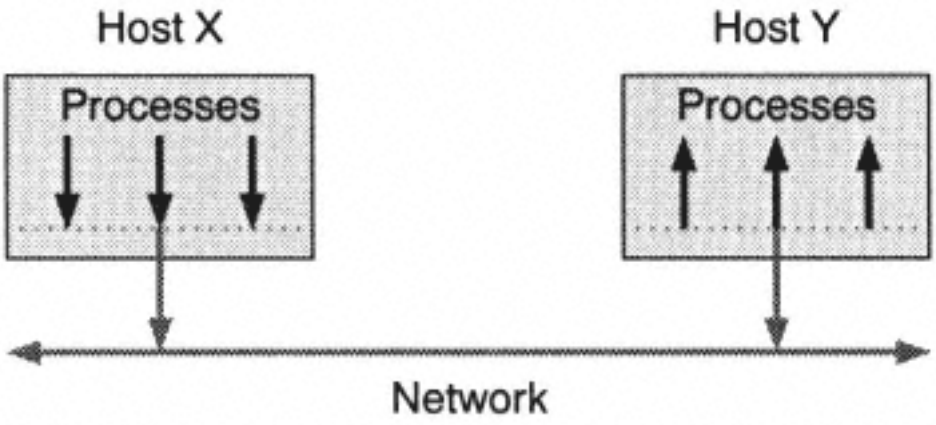

Figure 13.1
Communication software multiplexes and demultiplexes messages.

32-bit identifier is taken as the standard, but all identifiers are 16 bits long, space in the frame will be wasted whenever an identifier is transmitted.

From the two situations described above, it should be apparent that process identifiers are not necessarily the best method of identifying a process on a remote host. Ideally, a common process identifier structure is agreed upon by all operating systems; however, since agreement is unlikely, many communication systems have a second set of identifiers used for communications. Commonly known as ports, these identifiers have a structure that is agreed upon by all operating systems on the network. A process is associated with its operating-system-specific process identifier and is *bound* to as many ports as necessary for the process to perform its communication functions; processes not involved in network communications are not bound to a port.

For example, in the network shown in Figure 13.2, processes Py and Pz are bound to ports 2 and 3, respectively, on host E; on host S, processes P4, P9, and P1 are bound to ports 1, 2, and 4, respectively.

The host address and port number are often referred to as a single data structure known as a socket, which consists of a host-port pair. *Host* uniquely identifies the process's host, while *port* identifies the process's port. Should the process number change (for example, the host may crash and reboot, resulting in the software module obtaining a different process number), the system will still function as long as the software module gets the original socket.

Before a process can communicate, it must first bind to a socket. Thereafter, all messages that arrive on the host with the process's port number are returned to the process. When transmitting a message, a process supplies the communication system with the message and the socket of the destination process. The source communication system uses the host identifier to determine the destination host, while the destination communication system uses the port identifier to determine the destination process.

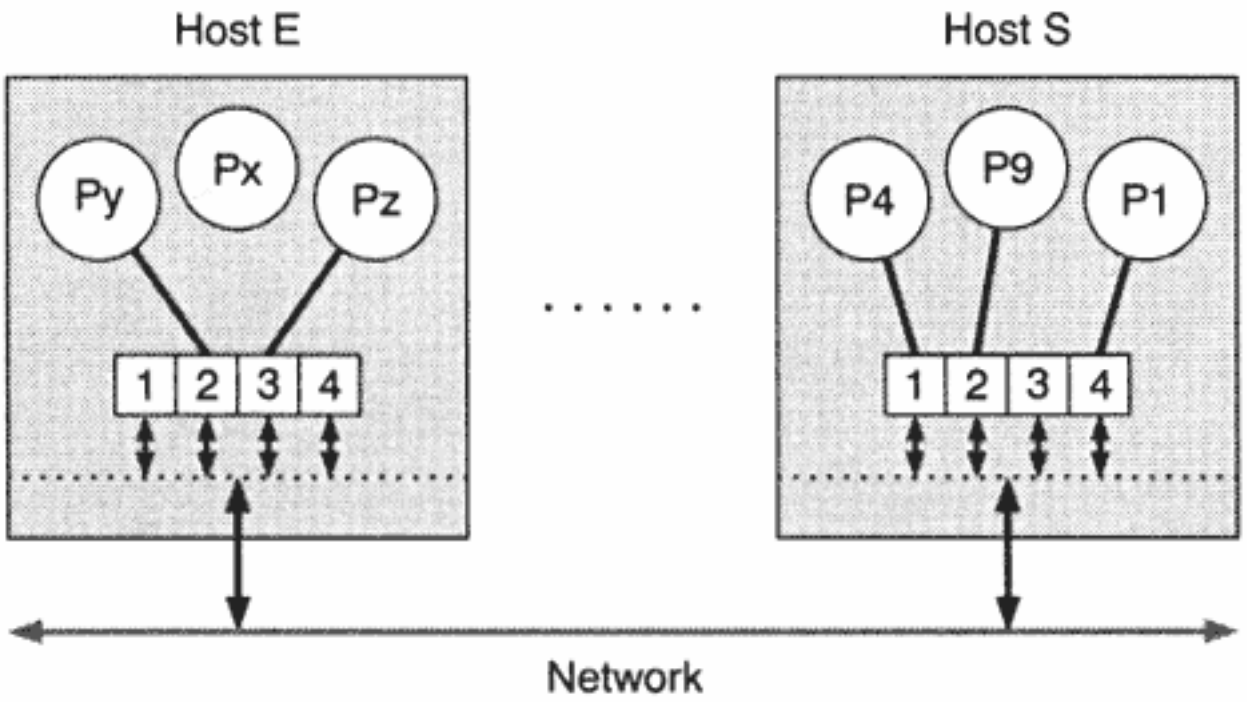

Figure 13.2
Processes are identified by their ports.

For example, using Figure 13.2, if process Py on Host E is to send a message to process P1 bound to port 4 on host S, the message is sent to socket address <S,4>.

Finally, it must be noted that sockets are not the only method of transmitting messages without using explicit process identifiers. For example, a mailbox is an intermediate entity that processes can write to and read from. If process A sends a message to process B, the message is written to process B's mailbox. When process B reads its mailbox, the message from A is returned.

13.2.2 Layering

The discussion in the previous section makes it clear that it is unrealistic to expect each process to support its own version of sockets and the communication software. Instead, most communication software has been developed to offer a range of communication services to processes and their related applications. Communication software is now typically written as a series of layers, each one offering services to the layer above it and using the services of the layer below it. There are many reasons for developing layered systems, including

Enhancing features. Layers can add features or facilities to a network. For example, an unreliable, error-prone physical layer can be made reliable through use of a data link layer supporting a go-back-N protocol.

Hiding features. The addition of a layer can also hide features. For example, it is possible to use names rather than network addresses when referring to a station; an application can supply a host name to an underlying layer, which maps the name into a network address.

Ease of modification. Layered software is typically easier to maintain than monolithic software, since the functions performed by a layer are localized to the module associated with the layer in question. Fault detection and subsequent module retesting are also simplified. For example, controlled module testing can be performed if the test software generates the necessary error conditions. This is typically faster than waiting for the error to occur while the module is running.

Portability. Layered software is typically more portable than unlayered software. For example, if an application is written for a specific network (say an 802.3 bus), transferring it to another network may well require extensive modifications. By writing application software as *network independent*. moving the application between networks can be a simple task.

Layering has been used extensively in Commkit; for example, the wide area network software consists of four distinct layers:

Physical. Responsible for the transmission and reception of bytes across the channel.

Data link. Defines the channel packet structure and is responsible for the orderly flow of information between the interconnected DTEs.

Network. Responsible for the establishment, maintenance, and eventual clearing of virtual circuits.

Application. In the wide area network, a network testing tool that permits the user to establish a number of virtual circuits with various DTEs.

13.3 Terminology

Each host attached to the underlying subnet is referred to as a system. Within a single subnet, all systems typically support the same layers; the layers themselves are said to form a layered architecture. The highest layer of a layered architecture consists of applications and processes using the communication services offered by the underlying layers. The lowest layer deals with the physical connection of the systems making up the subnet.

Although the layers perform different functions, there are three underlying similarities:

Each is associated with its own internal protocol. For example, in the Commkit wide area network, the data link layer supports a go-back-N protocol, while the network layer has a virtual circuit protocol.

Each offers services to the layer above it. As an example, the data link layer of Commkit's wide area network offers a reliable point-to-point communication service to the network layer.

Each uses the services of the layer below it. For example, the application software in the Commkit wide area network uses the virtual circuit service offered by the underlying network layer.

Because of the widespread use of layering, many layered architectures use a more formal description of each layer and its interactions with adjacent layers. For example, any layer *N* offers one or more N services to the layer above it, *N* + 1. Similarly, layer *N* uses the services offered by the layer below it, *N* 1. Layer *N* + 1 is referred to as the service user and layer *N* as the service provider.

Each layer *N* supports software (or hardware) modules known as protocol entities that conform to the N-protocol associated with the layer in question. Protocol entities that exist at the same layer are known as peer protocol entities; they exchange information in N-protocol data units, or NPDUs, using the services of the *N* 1-layer. The service user that receives information is referred to as the correspondent user. This is represented diagrammatically in Figure 13.3.

These definitions are recursive; that is, any layer *N* becomes layer *N* + 1

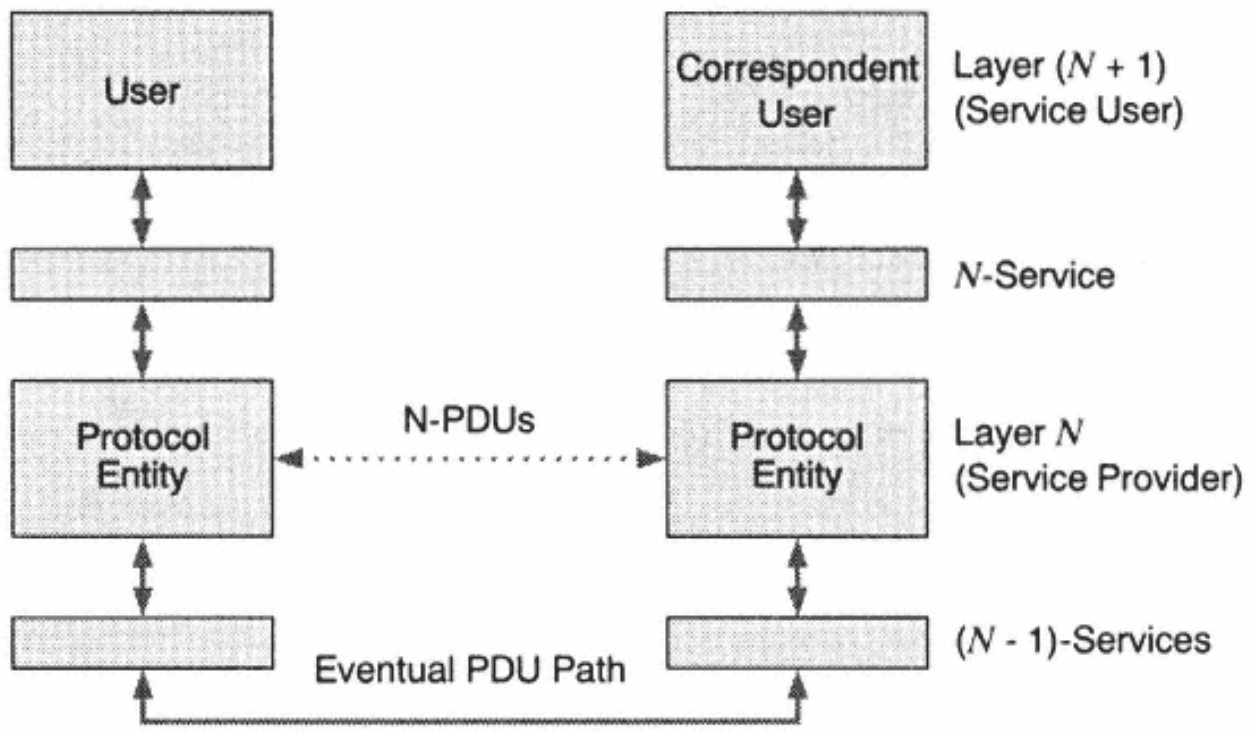

Figure 13.3
Terminology of a layered architecture.

the further into the architecture one proceeds. At the lowest layer, there is no layer *N* 1; the lowest layer must support the physical interconnection of the systems. Similarly, at the highest layer, there is no layer *N* + 1 other than end users or application processes.

Users at layer *N* + 1 access N-services through service access points, or SAPs. Since at any layer *N* + 1 there are potentially multiple users of a particular *N*-service, each service user is assigned a unique SAP address. The boundary between the two layers is known as an interface and is crossed through an SAP.

The *N*-layer PDU also contains *N*-layer peer protocol entity control information (such as sequence numbers and checksums) in the N-PCI (or N-protocol control information) as well as the protocol data unit from layer *N* + 1. At layer *N*. the protocol data unit from layer N + 1 is referred to as the N-SDU (or N-service data unit). The N-PDU *encapsulates* the (N+1)- PDU. The relationship between the *N* + 1 and *N* layers is represented in Figure 13.4.

The services offered by any layer fall into two categories: confirmed and unconfirmed (corresponding to the connection-oriented and connectionless services discussed in previous chapters). Support for the transfer of informa-

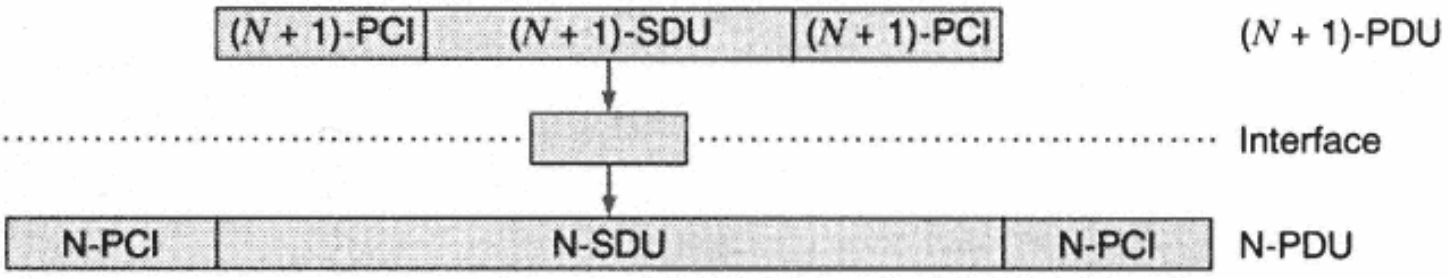

Figure 13.4
Encapsulation of the N-PDU.

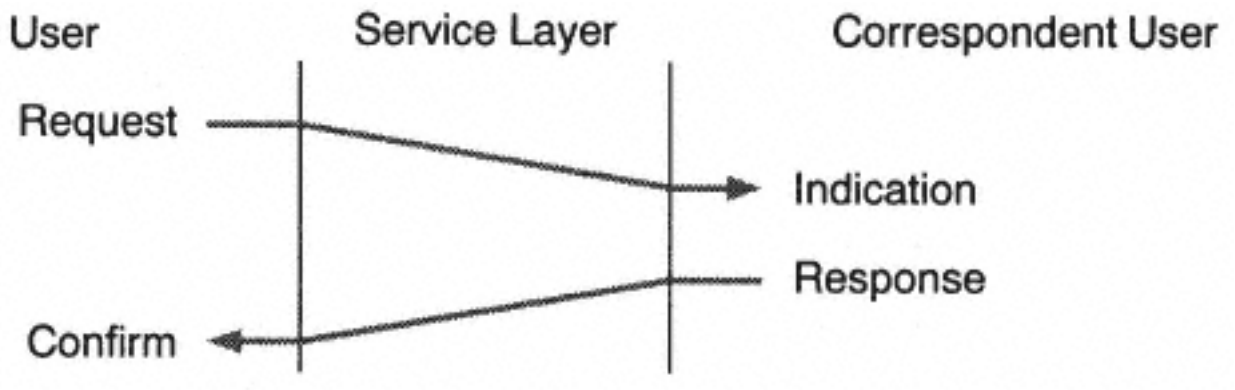

Figure 13.5
Information flow between a user and a correspondent user.

tion is described by four primitives (the first two are used by both confirmed and unconfirmed services, while the last two are for confirmed services only):

Request. A service user initiates a transfer to a correspondent user by passing a request primitive to the service provider. The service provider takes this information (an SDU) and encapsulates it into a PDU. The PDU is sent to the correspondent service provider on the remote system using the services provided by the underlying layers.

Indication. Upon reception of the PDU, the correspondent service provider supplies the correspondent user with an indication primitive.

Response. The correspondent user is expected to acknowledge receipt of the indication with a response primitive. The correspondent service provider returns the response to the originating system in a PDU.

Confirm. When the original service provider receives the PDU, a confirm primitive is returned to the user.

The information flow between the (N + 1)-layer users and the N-layer service is shown in Figure 13.5 (again, the response and confirm primitives are used only in confirmed services).

Many layered architectures have adopted the `layer.service.primitive` notation when referring to the services required by a user at a particular layer. Typical services include CONNECT (for establishing a connection), DATA (for the transfer of data), DISCONNECT or CLEAR (for clearing a connection), and RESET (for resetting a connection). For example, when establishing a virtual circuit in the Commkit wide area network, the application issues a CONNECT request, written as `A.CONNECT.request` and including the address of the destination application.

13.4
Examples of Layered Architectures

Layered architectures are by no means a new phenomenon in data communications. Both IBM's System Network Architecture (SNA) and Digital's Digital Network Architecture (DNA) were released in the mid-1970s. There are several other well-known architectures, three of which are examined in this section: the Open Systems Interconnection model, IEEE 802, and the DARPA Protocol Architecture.

13.4.1
The Open Systems Interconnection Reference Model

Proprietary products such as SNA restrict users to a specific manufacturer's line of equipment. When SNA was introduced, customers and manufacturers alike recognized this problem and in the late 1970s began pushing for an "open" architecture for which all manufacturers could develop products. The objective of such an architecture was to allow customers with equipment from different manufacturers to communicate on the same subnet, given adherence to a set of agreed-upon standards.

Probably the best known example of an open architecture is the Open Systems Interconnection (or OSI) reference model proposed in the late 1970s by the International Organization for Standardization (ISO). The OSI model consists of seven layers, each of which supports one or more services. Although a standard, OSI is not an implementation standard; that is, the model explains what the seven layers should do, but does not specify the services and protocols associated with each layer. (Section 13.3 described layering using OSI terminology.)

The main objective of the OSI model is to support communications between end users on different systems. The end users are not part of the OSI model, they simply use the facilities it offers. The four uppermost layers (application, presentation, session, and transport) support end-to-end protocols that are network independent. The three lowest layers (network, data link, and physical) support protocols that are network dependent.

The diagram shown in Figure 13.6 illustrates the logical and physical flow of information between a pair of end users on different systems. Messages exchanged between end users flow vertically through the system until they reach the Physical Layer, at which point they are transmitted between systems ("xPDU" denotes the PDU used by a specific layer; Packets, Frames, and Bits normally are not referred to in terms of PDUs).

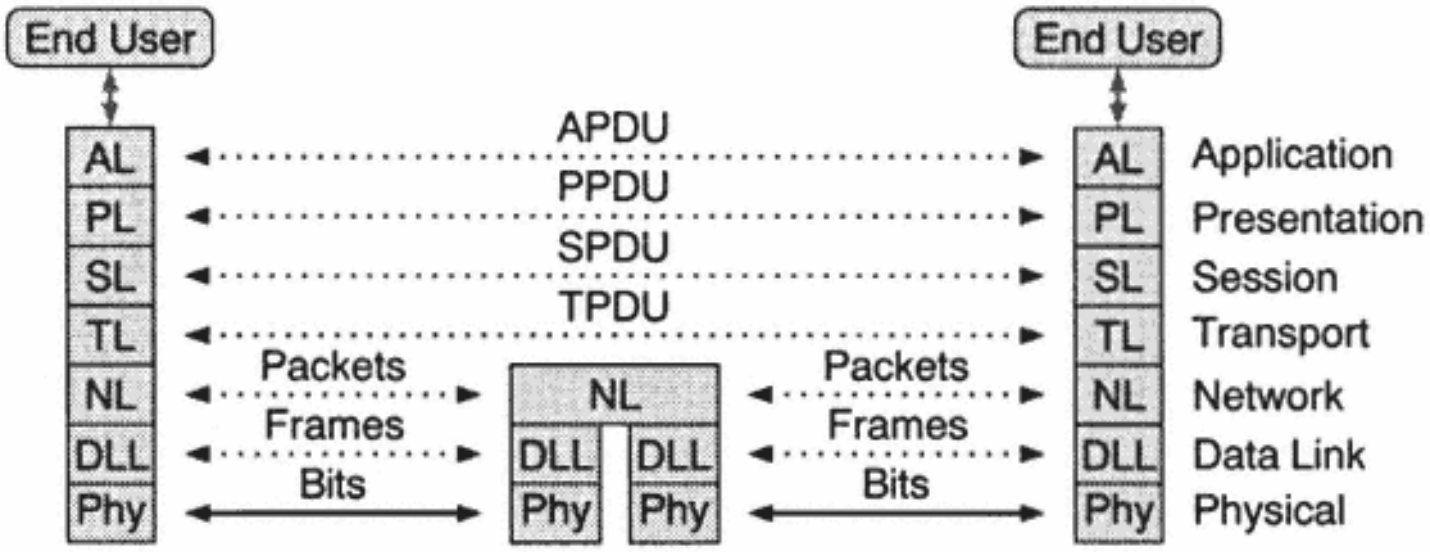

Figure 13.6
The Open Systems Interconnection reference model.

Network-Dependent Layers

The lowest three layers of the OSI model (physical, data link, and network) are network-dependent, defining protocols for the transfer of information between systems.

The physical layer is the lowest layer in the OSI model (layer 1), responsible for the transmission of bits over a physical channel; it covers issues such as signaling, modulation, and defining how components are physically interconnected. Examples of standards for the physical layer include RS-232-C, RS-449, and X.21. Physical layer standards for local area networks include IEEE 802.3 (ISO 8802.3), IEEE 802.4 (ISO 8802.4), and IEEE 802.5 (ISO 8802.5). Newer physical layer standards include those for ISDN networks, such as I430/1.

The second layer in the OSI model is the data link layer, responsible for the transfer of information in the form of frames across the underlying physical layer. The data link layer can be reliable, overcoming physical layer errors using go-back-N protocols. Examples of data link layer protocols include HDLC, SDLC, ADCCP, and LAP-B. In some networks, the data link layer is unreliable and each frame is treated as a datagram. Examples of unreliable data link layer protocols are the LAN protocols 802.3 and 802.5.

The network layer is the third layer of the OSI model, handling the routing of packets across the underlying subnet. In connection-oriented networks, it is responsible for call establishment, data transfer, and call clearing. ITU X.25 is an example of a network layer protocol. ISO standards include ISO 8473 and 8348.

Network-Independent Layers

The lower three layers of the OSI model describe point-to-point communications between pairs of nodes on a subnet. The remaining layers assume that the underlying layers can carry information to a given remote system.

The fourth layer, the transport layer, defines how end-to-end communications can be established and maintained across the subnet. One of its functions is to make the network transparent to the remaining upper layers of the OSI model. OSI defines five classes of connection-oriented transport service. Each class offers a specific quality of service based upon the underlying network layer's quality of service:

Simple class (protocol class 0). The simple class operates over networks with low failure and error rates. A class 0 transport layer service therefore does not enhance the service offered by the network layer, leaving to the network layer sequencing and flow control. In class 0, there is one network connection for each transport connection; should the network connection fail, the transport connection will also fail.

Basic error recovery class (protocol class 1). Protocol class 1 is similar to class 0, with the exception that it operates on networks that are subject

to network resets. For reset recovery, it is necessary for the transport service to maintain TPDU (Transport Protocol Data Unit) sequence numbering.

Multiplexing class (protocol class 2). A class 2 transport service allows the multiplexing of one or more transport connections over a single network connection. The underlying network is assumed to be fully reliable.

Error recovery class (protocol class 3). Protocol class 3 supports the features found in classes 1 and 2-notably, it can recover from network resets, and it supports multiplexing.

Error detection and recovery class (protocol class 4). A class 4 network operates atop a datagram network, meaning that extensive error detection and error-handling features are necessary (i.e., sequence numbering, CRC checking, timeouts, and TPDU retransmission).

Transport layer protocols include ISO 8072 (OSI transport service), ISO 8073 (OSI transport protocols), and ITU X.214 and X.224.

The fifth layer of the OSI model is the session layer. Unlike the lower layers, which are concerned with data transportation, the session layer deals with the management of interactions (or dialogs) between the two end users. Each session connection has a corresponding transport connection. Should that connection fail, the session layer can re-establish the connection transparently (i.e., without end-user intervention). Similarly, once an end user has completed a session, the session layer may choose to keep the transport connection active, initiating another session across the same one.

The session layer may offer a number of data transfer services. It supports both full- and half-duplex user dialogs. In half-duplex, the transmitting user possesses a permission-to-transmit token that can be exchanged by the users. Other data transfer services include quarantining and synchronization. Quarantining allows the local session layer to withhold a number of messages until explicitly instructed to deliver them by the remote session layer. Synchronization allows a data transfer to "rewind" to a checkpoint specified by the end-user application. Both the quarantining and synchronization services are useful in transaction-oriented applications where the loss or retransmission of information may lead to inconsistencies in, for example, a database.

In a truly open system, there is no guarantee that the equipment available to one end user will be the same as that available to another. If the two end users are to communicate, it may be necessary to convert the information format used on one system to that used on another. For example, one system may use ASCII and the other EBCDIC; similarly, the internal representation of integer or floating-point numbers may differ, requiring some form of conversion. Layer 6 of the OSI model, the presentation layer, deals with the representation and transformation of the data between application layer entities.

In its simplest form, each system's presentation layer can support a translator for every other possible system on the network. If there are *N* systems on the network, a total of $N \times (N\ 1)$ translators are needed. However, if a common network-wide information representation is chosen, then only $2 \times N$ translators are needed: one from the machine's internal representation to the network external representation, and the other from the external representation to the internal representation.

Presentation layer protocols include ITU X.409, Presentation Transfer Syntax, and GM's MAP Standard Message Format. In GM's presentation layer protocol, internal data is converted into a structure consisting of a type identifier (for example, Boolean, Integer, or Floating Point), an optional length indicator (depending upon the type), and the converted value. Booleans are a single byte (`0x00` for true and `0xFF` for false), while a floating-point number is converted into an ASCII string.

The presentation layer can also be responsible for the encryption and decryption of information, as well as for file compression.

The application layer is the uppermost layer of the OSI reference model, offering application services to the different end users. Some of the more common application services include file transfer protocols, electronic mail (such as X.400), and virtual terminals (such as X.28 and X.29).

To avoid the duplication of work at the application layer, three types of application elements are proposed by the OSI model for use by end-user applications:

Elements within the end-user application that deal with accessing OSI services (User Elements, or UE).

Elements that have capabilities useful to a variety of applications (Common Application Service Elements, or CASE).

Elements that have capabilities for specified applications (Specific Application Service Elements, or SASE).

13.4.2 IEEE 802

The lower three layers of the Open Systems Interconnection model are for connection-oriented wide area networks. These WANs exhibit a number of common properties, including low speeds (typical maximums of about 56 kbps), high error rates (about 1 bit in 105), and the need to perform packet routing from node to node. Accordingly, the protocols proposed and developed for the lower levels of the OSI model overcome these limitations.

Since local area networks such as the Ethernet and the IBM Token Ring are faster (speeds of 10 Mbps and up), exhibit lower error rates (about 1 bit in 109), and do not require routing, many of the functions provided by the lower OSI layers are redundant or unnecessary in a LAN environment. In short, many of the protocols associated with the lower layers of the OSI model are considered to be too *heavyweight* for LANs.

The IEEE recognized these limitations and set about defining the 802 standard, intended for high-speed, low-error-rate local area networks. 802 protocols are *lightweight*. As with the OSI model, one of the objectives of the 802 standard is to ensure that equipment from different manufacturers can communicate if the standard is adhered to.

The 802 standard is a three-layer architecture that encompasses the physical and data link layers of the OSI model. Figure 13.7 shows the overall structure of the 802 protocol family and its relationship to the OSI model.

The IEEE-802 standard consists of the following parts:

802.1 Overview, internetworking, and systems management; defining the relationship between the various 802 protocols and the ISO higher-layer protocols.

802.2 Logical Link Control (LLC) sublayer.

802.3 CSMA/CD bus access method (described in Chapter 11).

802.4 Token-passing bus access method (described in Chapter 12).

802.5 Token-passing ring access method (described in Chapter 12).

802.6 Metropolitan Area Network (MAN) access method (described in Chapter 12).

The various 802 physical layers were described in Chapter 11 (802.3 CSMA/CD) and Chapter 12 (802.4 Token Bus, 802.5 Token Ring, and 802.6 MAN). Additional 802 physical layer protocols are currently under development, including 802.3u (100 Mb/s CSMA) and 802.12 (100 Mb/s Token Ring).

The OSI data link layer is divided into two parts: the MAC (Media Access Control) sublayer and the LLC (Logical Link Control) sublayer.

Media Access Control

The MAC sublayer provides three information exchange primitives to the LLC sublayer, regardless of the underlying network. These primitives are as follows (note that OSI terminology is used):

`MA.DATA.request`. Takes three parameters from the LLC sublayer: the network address of the destination SAP (either a unicast, broadcast, or

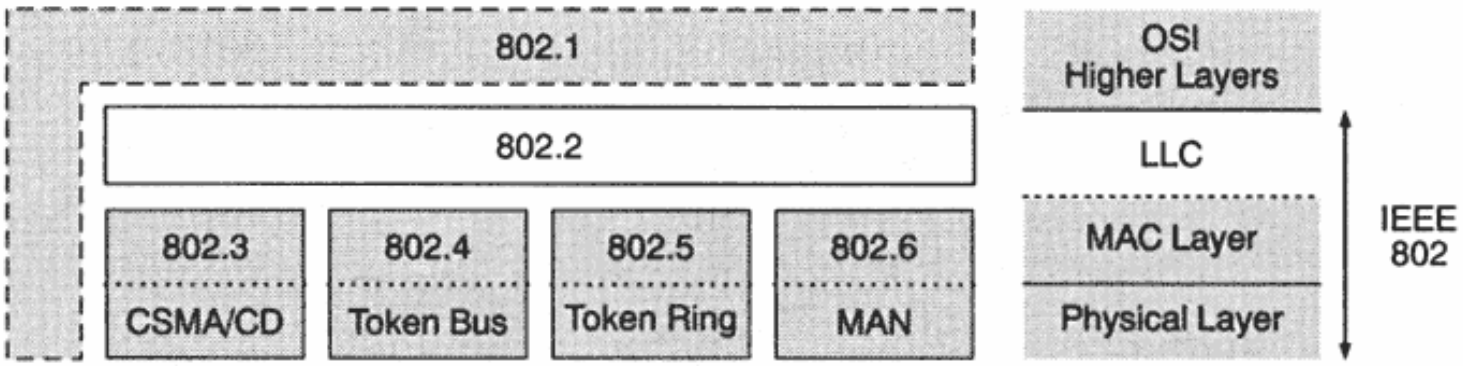

Figure 13.7
The 3-layer 802 protocol family.

multicast address), the SDU to be transmitted (an LLC Protocol Data Unit), and a service class requesting a certain priority level. The SDU is transmitted by the MAC layer using the physical layer services of the underlying network.

The service class is used by networks that can support different levels of priority, such as the 802.5 Token Ring. The MAC sublayer transmits the frame with the specified priority. Networks such as 802.3 that do not support priorities simply ignore the service class.

`MA.DATA.indication`. Returns the source and destination address, as well as the SDU received, to the correspondent LLC sublayer. The status of the reception is also returned, indicating either that the frame was received correctly or the reason for the failure of the reception.

`MA.DATA.confirm`. Returns the status of the `MA.DATA.request` to the LLC user, verifying the completion of the transmission or specifying a reason for the failure. The result returned indicates only the success of the transmission at the MAC sublayer. In networks such as the 802.5 Token Ring, the confirmation is extracted from the status bits in the ring frame. However, in the 802.3 CSMA/CD network, the confirmation is generated locally; an example of an 802.3 transmission failure is when the collision limit is reached.

A time-sequence diagram of the primitives described above and the MAC and LLC sublayers is given in Figure 13.8.

Logical Link Control

The Logical Link Control (LLC) sublayer offers three types of service to the upper OSI layers: unacknowledged connectionless; connection-oriented; and acknowledged connectionless. The connectionless services are for applications that support their own error-checking and recovery schemes.

The unacknowledged connectionless service supports two service primitives: `L.DATA.request`, for the transfer of an SDU to a given remote address; and `L.DATA.indication`, signaling the receipt of an SDU. This is a datagram service; there are no acknowledgments associated or supported with it. The time-sequence diagram for the unacknowledged connectionless service is shown in Figure 13.9.

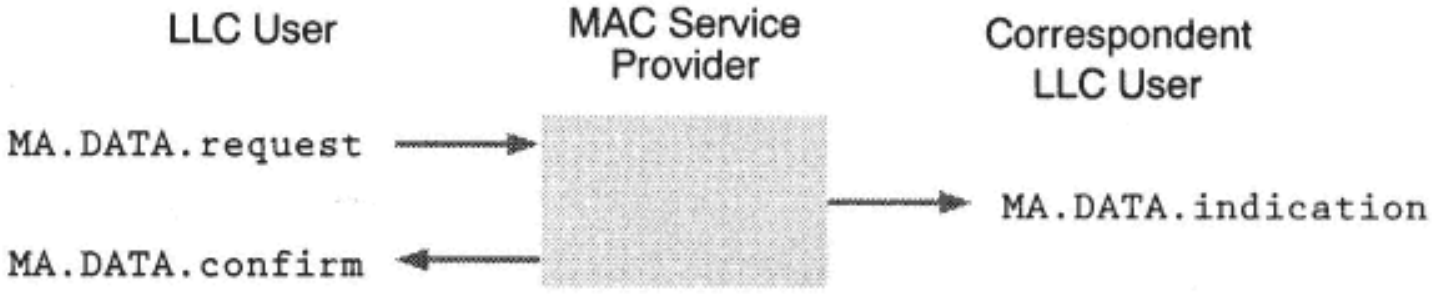

Figure 13.8
MAC layer information interchange primitives.

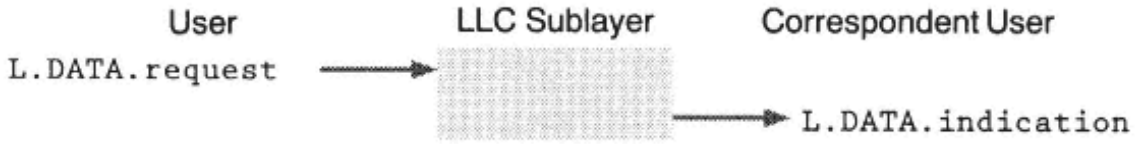

Figure 13.9
LLC Unacknowledged Connectionless Service.

In the acknowledged connectionless service, an `L.DATA_ACK.request` contains an SDU, the delivery of which is acknowledged by the remote LLC sublayer, as shown in Figure 13.10.

The connection-oriented service shown in Figure 13.11 allows the service user to request the establishment of a connection using `L.CONNECT.request` (note that there is no `L.CONNECT.response`, since the response is generated by the correspondent LLC sublayer).

Once the connection has been established, either service user may initiate a data transfer using `L.DATA_CONNECT.request` (Figure 13.12). The connection can be cleared by either user issuing an `L.DISCONNECT.request` (Figure 13.13).

In all three types of service, the addresses supplied to the LLC SAP

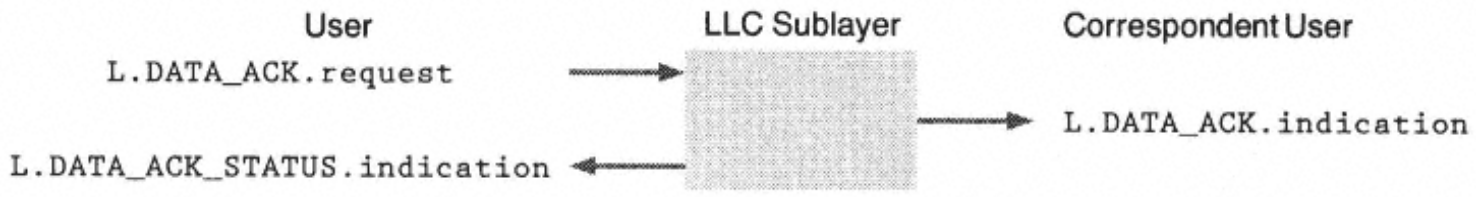

Figure 13.10
LLC Acknowledged Connectionless Service.

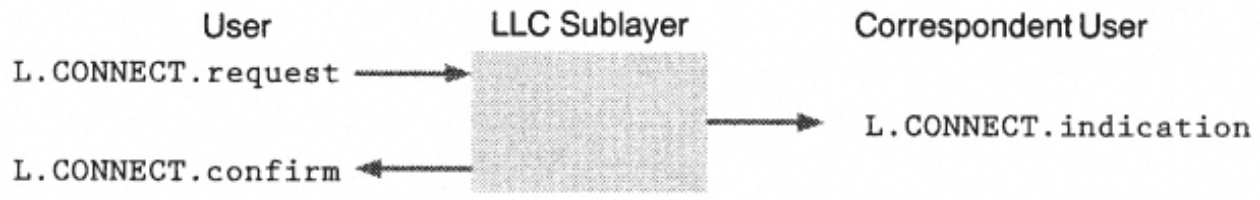

Figure 13.11
LLC Connection-Oriented Protocol: Connection Request.

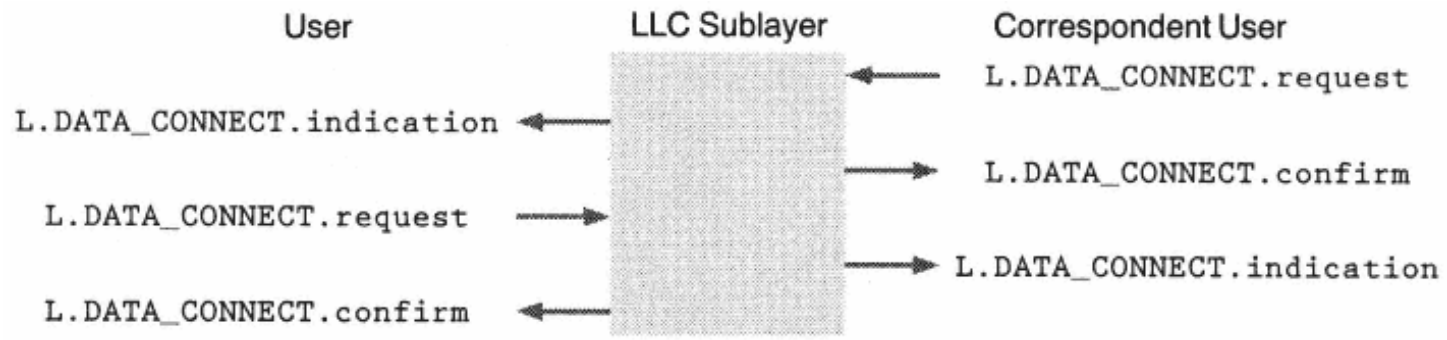

Figure 13.12
LLC Connection-Oriented Service: Data Transfer.

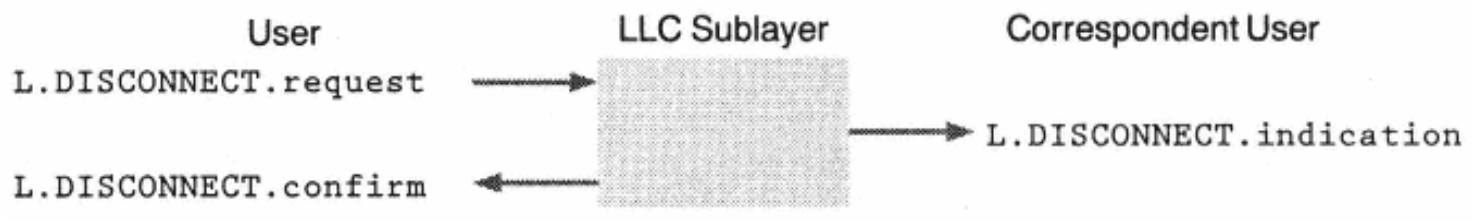

Figure 13.13
LLC Connection-Oriented Service: Call Clearing.

Figure 13.14
The LLC Protocol Data Unit (all sizes in bytes).

indicate the destination LLC SAP as well as the destination physical address. Regardless of the service type, there is a single LLC Protocol Data Unit (LLC PDU) format, as shown in Figure 13.14. The LLC PDU fields are described as follows:

DST SAP. The destination LLC service access point address. One bit indicates whether the SAP is an individual or group address, meaning that any one of 128 individual or 128 group LLC SAPs may be specified.

SRC SAP. The LLC service access point for the source of the Information.

Control. Each LLC PDU is sent with a control field, the format and contents of which are based upon the HDLC protocol control field (see Chapters 6 and 10). The connection-oriented service supports both Information and Supervisory control fields. On the other hand, the connectionless services use the Unnumbered Information (UI) control field.

Information. A variable-length data field, *M* is a multiple of 8 bits. It can be 0for example, when sending supervisory PDUs.

There is no LLC checksum. The MAC layer FCS is presumed to offer sufficient error detection capabilities.

13.4.3
The DARPA Protocol Architecture

The DARPA protocol architecture (or DPA) is an open system architecture built for a single (large) customer, the U.S. military. It was originally developed in the mid-1970s as part of the military's DARPA (Defense Advanced Research Project Agency) study of internetwork communications. Earlier research funding from DARPA (or ARPA as it was known in the 1960s) went into the development of the ARPANET packet-switching network.

The DARPA protocol architecture consists of three network-independent layers, built atop the network access services of any underlying network. The layers and their relationship to the OSI model are shown in Figure 13.15.

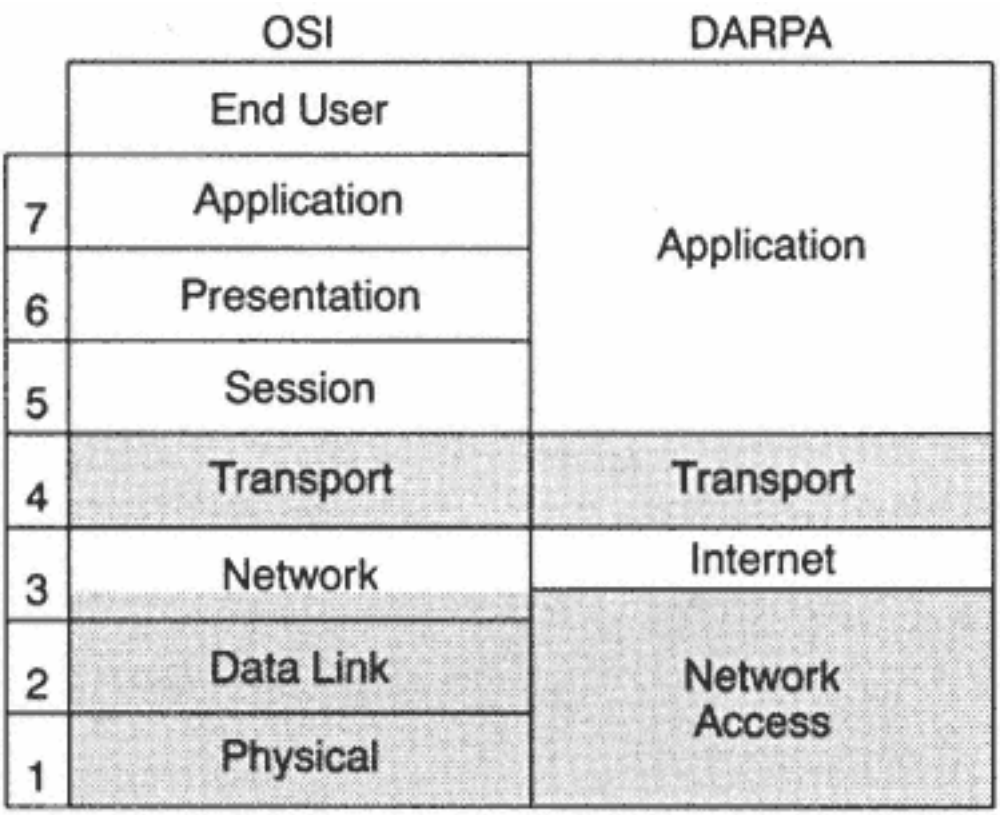

Figure 13.15
The relation between the OSI model and the DARPA protocol architecture.

The DPA operates atop almost any type of network, as long as the necessary software exists between the Internet and Network Access layers. The underlying network can be connection oriented or connectionless, a wide area or local area network. The type of service offered by the underlying network is transparent to the top three layers.

In order to achieve this transparency, the designers of the DPA had to adopt several standards. For example, since byte ordering of an integer can vary from machine to machine (see Chapter 1), machines have been classified as either little endian (in which the low-order byte is associated with the lowest memory location) or big endian (where the low-order byte is stored in the highest memory location). The standard adopted by DPA requires all machines to transmit information in big endian style (i.e., the most significant byte first).

The need for network-independent host addresses was also recognized. Since different networks support different address structures, each host is assigned a unique 32-bit Internet address. When a packet is available for transmission, it is supplied with the destination host's Internet address in an Internet datagram to the Internet layer. The Internet layer maps the Internet address into a network-specific physical address; this address is then used by the Network Access layer as the destination address.

Should the Internet layer not recognize the Internet address, an Address Resolution Protocol (ARP) packet containing the unknown Internet address is broadcast on the network. The station with the address in question returns its physical address in an ARP reply packet.

The Internet layer offers a connectionless datagram service to the Transport layer, known as the Internet Protocol (IP). The Internet datagram contains two parts: a header and a data area. The Internet layer can fragment

datagrams if they are too large for the underlying network. The datagram header contains control information as well as the destination and source Internet addresses.

The Transport layer offers two levels of service to applications:

A datagram protocol known as UDP, or User Datagram Protocol. UDP messages are supplied to the Internet layer for transmission to a given Internet destination. Since a single host may have many processes transmitting and receiving UDP messages, each host has a number of UDP ports to which the messages are addressed.

Upon receipt of a UDP message from the underlying network, the Internet layer supplies the message to the UDP service, which makes the message available to the process associated with the port specified in the message.

A reliable, full-duplex, connection-oriented stream protocol known as TCP, or Transmission Control Protocol. A message sent by TCP is sent in a segment consisting of a TCP header and a data part. The header contains the source and destination ports as well as a 32-bit sequence number and a 32-bit acknowledgment number. TCP is a sliding window protocolthe sequence and acknowledgment numbers specify the current byte number and the number of the next expected byte, respectively. The window size is expressed in bytes, and each segment is sent with a checksum. TCP is usually referred to as TCP/IP, since the TCP protocol is interwoven with IP.

Operating systems such as UNIX 4.4BSD support both TCP/IP and UDP. In UNIX 4.4, the combination of the host Internet address and the port number is known as a socket.

The aforementioned three layers of the DPA are commonly referred to as the Internet. Most users deal only with the highest layer of the DPA, the Application layer, which consists of cooperating processes or tasks transferring data via TCP/IP and UDP. Most Internet applications are described in terms of the client-server model (originally described in Chapter 2), in which a process offering a service is the server, while the process requiring the service is the client. In the client-server model, the client sends a request to the server; upon receipt of the request, the server performs the service. When finished, the server sends a response to the client. Since a process on a host is associated with a port, many services are associated with "well-known" ports, which facilitates the identification of the service by a client.

Internet application services include

SMTP. Simple Mail Transfer Protocol, the Internet electronic mail protocol. SMTP operates atop TCP/IP and defines the steps in transferring electronic mail between two hosts.

Rlogin. A remote login service offered by UNIX 4.4BSD, allowing users on a UNIX machine to connect to another, remote UNIX machine. Since

multiple remote logins are possible, each login, once established, communicates through its own TCP/IP port.

TELNET. Similar to rlogin, except that TELNET is intended for any remote Internet host. It operates atop TCP/IP.

FTP. File Transfer Protocol, a reliable file transfer protocol that uses both TCP/IP and TELNET. TCP/IP handles the transfer of the file; TELNET allows the user to sign on to the remote host from which the file transfer is to take place.

HTTP. HyperText Transfer Protocol, the most recent development of the World Wide Web, a tool intended to encourage and simplify access to the Internet. An HTTP client allows a user to access and display hypertext documents with links to other documents; an HTTP server returns hypertext documents.

In all of the above examples, each service is associated with its own well-known port; for example, SMTP is associated with port 25, TELNET with port 23, and HTTP with port 80.

The OSI model, unlike DPA, offers reliability at each of the four lowest layers. In the DPA approach, the application service handles reliability on an end-to-end basis.

Although the DARPA Protocol Architecture is used widely for intranetwork communications, it is actually intended for internetwork communications. The internetworking aspects of the Internet protocol are discussed in more detail in Chapter 14.

13.5 Commkit Sockettes

Commkit supports a simplified version of UNIX UDP sockets, known as sockettes, which offer the foreground process a number of communication primitives that hide the underlying network. The sockette software is designed to work with any of the Commkit networks (wide area, bus, and ring). Other than a few lines of specialized code for the calls to individual networks, the sockette software is identical for all networks, meaning that the replacement of one network by another is transparent to the user software.

Each host in a sockette network is assigned a name and a corresponding internal sockette host number. This information is stored in the file `hosts`. The physical station addresses are stored in a separate file, `addrs`, and contain the internal sockette host number along with the associated physical address.

Each foreground process can be associated with up to three sockettes. The sockette itself consists of two parts: a host number and a port number. Host numbers are obtained by the process either from an incoming sockette message or by supplying a host name to the `gethostbyname()` primitive

(see below). There are ten different port numbers, some of which are "well known" and intended for specialized applications. For example, `ECHO_PORT` is port 2 and is used by a foreground process that echoes each message that it receives.

13.5.1
Sockette Data Structures

The sockette data structures are defined in `sockette.h` and `sockcode.h`. Those in `sockette.h` are specific to the primitives found in `sockette.c`, while those in `sockcode.h` are sockette return codes and other data structures intended for sockette application software.

A sockette address (defined in structure `sockaddr`) consists of two fields, a port address, `port` (of type `port_address`, an unsigned character) and a host address, `addr` (of type `host_address`, an unsigned integer). A process *must* be bound to a sockette address to participate in a communication.

Regardless of the network, all sockette messages are sent as part of the `network_msg` structure (Figure 13.16).

The fields of the sockette message are defined as follows:

`net_dst`. The physical network address of the destination host of the message. This is obtained from file `addrs`.

`net_src`. The physical network address of the source host of the message. This is obtained from file `addrs`.

`sock_dst`. The sockette address of the destination process, consisting of the destination host and port numbers. `Sock_dst` is generated by the transmitting process.

`sock_src`. The sockette address of the process sending the message; contains the host and port numbers of the transmitting process.

`net_msg`. An array of up to `MAX_NET_SIZE` bytes of data.

The network message is transmitted in the underlying network's frame according to the rules associated with the network in question.

The file `hosts` consists of one or more host numberhost name pairs, entered on separate lines. The host number is any positive integer, while the host name is any name up to nine characters in length; the host number is separated from the host name by one or more spaces. The `hosts` file is read once, into the data structure `host_addrs`. The maximum size of `host_addrs` is set in `NUM_HOSTS`.

Figure 13.16
The sockette message structure.

Addresses are kept in the file `addrs`, in which each entry is a host number and the physical address of the station, separated by a single space. The host number is a positive integer, and, in keeping with Commkit's network address structure, the physical address is a single character. `Addrs` is also read only once, into the structure `addr_to_physical`, the maximum size of which is set in `NUM_HOSTS`.

13.5.2
Sockette Primitives

The Commkit sockette software supports seven primitives, all of which are found in `sockette.c`:

`int gethostbyname(name, *address)` scans the list of host names to find a name that matches the name supplied (a character string). If name exists, the address of the name is returned in `address` (a pointer to a `host_address` structure). If the name is found, `SUCCESS` is returned; otherwise, `FAIL`.

`int getport(*portnum)` returns the *next* available port in `portnum`. An explicit port can be requested by entering a positive, non-zero value in `portnum` (a pointer to a structure of type `port_address`). If no ports remain or the requested port is already in use, `FAIL` is returned; otherwise, `SUCCESS`.

`int bind (*address)` attempts to bind the process with the supplied sockette address, `address` (a pointer to a `sockaddr` structure). It returns one of the following codes:

If the port specified in the `address` is invalid, `bind()` returns `FAIL`.

If the sockette's port address is already in use or all of the sockettes allocated to the process are in use, `bind()` fails, returning an error code of `PORT_IN_USE`.

If the sockette address is acceptable, `bind()` returns a sockette number, which has a value between 0 and `TOTAL_SOCKETTES` and must be used by the process whenever the address is to be used as the source address of a message.

`int release(sock_no)` attempts to dissociate the process from the specified sockette number, `sock_no`. The three possible return codes from `release()` are

`SUCCESS`returned if the `sock_no` supplied by the process is active.

`FAIL`returned if the value of `sock_no` is outside the range of legal sockette numbers.

`NOT_BOUND`returned if the process is not associated with the supplied sockette number.

`int get_net_address(sock_ host_addr, *phy_host_addr)` converts the sockette host address pointed to by `sock_host_addr` to a

physical host address, a network address pointed to by `phy_host_addr`. If the address doesn't exist, `FAIL` is returned; otherwise, `SUCCESS`.

`int sendto(sock_no, *buf, len, *to)` sends the message pointed to by `buf` of length `len` to the host-port pair specified by the socket address, `to`. If the size of the message is greater than the allowable message length (specified in `MAX_NET_SIZE`) or the process attempting to transmit isn't bound to the sockette number, `sockette, sendto()` fails and returns an error code of `BAD_SIZE` or `NOT_BOUND`, respectively. The source socket is obtained by mapping the socket number `sock_no` into the address supplied to `bind()`.

Note that since the to address can specify any process bound to any sockette (on any machine), it is possible to transmit messages between sockettes on the same machine.

`int recvfrom(sock_no, *buf, len, *from)` returns the first available message received for the process's sockette number `sock_no` (i.e., the address to which the sockette number refers). The message is copied into the memory pointed to by `buf`. The sockette address of the transmitting process is extracted from the network message and placed into `from`. The size of the message is returned if a message is available.

If `sock_no` is outside the range of permissible sockette numbers, `FAIL` is returned. A length greater than `MAX_NET_SIZE` results in a return code of `BAD_SIZE`. If no message is available, `FAIL` is returned (i.e., `recvfrom ()` supports asynchronous reception only).

In addition to the above, `sockette_init()` must be called once, prior to any sockette activity. `Sockette_init ()` initializes the various sockette tables.

13.5.3
Sockette Software Layering

The five layers of the Commkit sockette software are shown in Figure 13.17 and defined as follows:

Foreground Process. All applications using the sockette software must be written as the foreground process (the reason for this is explained later). For communications to take place, the process must first bind to a sockette; thereafter, messages can be sent.

Sockette Layer. The sockette layer is called by the foreground process whenever a sockette operation is required.

Network Layer. The network layer determines the destination of any message received (either from the network emulation software or from the foreground process). All messages are `network_msgs` that are received and forwarded using the `recv()` and `send()` primitives, respectively.

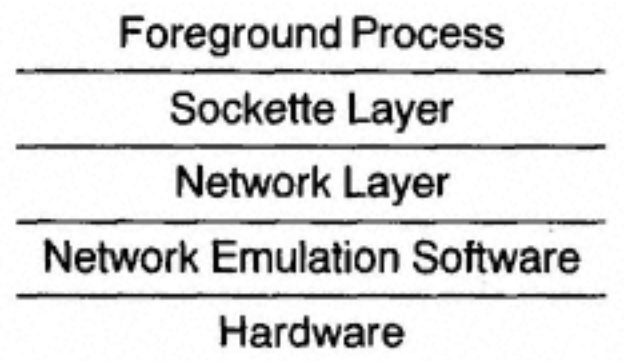

Figure 13.17
The sockette software layers.

Network Emulation Software. The network emulators remain unchanged; all messages received are forwarded to the network layer through queue `BACKGROUND_1`.

Hardware. The hardware interrupt handlers remain unchanged.

13.5.4
Sockette Implementation

The Commkit sockette software is designed to allow sockette communication between foreground processes. It is divided into two parts: a high-level part consisting of the sockette primitives described in Section 13.5.2; and a low-level part, handling sockette messages from the network and the foreground process.

At any moment, the foreground process can be bound to as many as three sockettes; each active sockette is assigned a sockette number that the process uses when sending and receiving sockette messages. A sockette number is returned after a successful `bind()` has taken place. This number is an index into the data structure `proc_info`, an array indicating the status of each sockette (i.e., whether the sockette is in use) and, if the sockette is in use, the sockette's address, which consists of a host-port pair. Whenever a `sendto()` takes place, the sockette number is used as an index into `proc_info` to obtain the sockette address; this address is the source address of the sockette message.

The mapping of textual host names into physical addresses is a two-step operation: first, the array `host_addrs` maps the name into the sockette host identifier. Then the sockette host identifier is mapped into a station's physical address by searching the array `addr_to_physical`.

Each host supports a total of `TOTAL_PORTS` ports, and the foreground process can bind to any port that is free. The data structure `port_tbl` contains the status of each port: whether it is free or has been bound to and, if bound, the queue associated with it. When the foreground process calls the `recvfrom()` primitive, the port associated with the sockette (obtained by indexing `proc_info` and extracting the port number) is used to index `port_tbl`. If a message is available on the port's queue, it is returned to the foreground process.

Regardless of the source of a sockette message (i.e., either the local foreground process or the network from a remote process), all sockette messages are sent to `network_layer()`, which compares the physical destination host address in the network packet with the host's address (`stn_id`). If these two are identical, the destination port number is used to determine whether the port has been bound to. The message is put on the destination port's queue if the foreground process is bound to the port in question.

If the packet is not destined for this host, `network_layer()` attempts to forward it. The forwarding algorithm is network dependent and selected at compile time in `makefile`:

Wide Area Network. All messages (from either serial port or the foreground process) are sent to `network_layer()`, which routes the message by examining the host's identifier and the message source (i.e., a serial port or the foreground process). Messages to unknown destinations are discarded.

Ring Network. The source of all messages is the foreground process; the message is sent to queue `RINGNET` for subsequent transmission.

Bus Network. The source of all messages is the foreground process; the message is sent to queue `BUSNET` for subsequent transmission.

The overall flow of messages is presented in Figure 13.18.

The design and final implementation of the sockette software were influenced by the following issues:

1. Since a special initialization call to the network layer process is necessary to bind to a sockette, the decision was made to allow only the Commkit foreground process to use sockettes.

2. If sockette messages are forwarded by the network layer to the foreground process, the sockette primitive `recvfrom()` is not necessary because the message is obtained using a `recv()`.

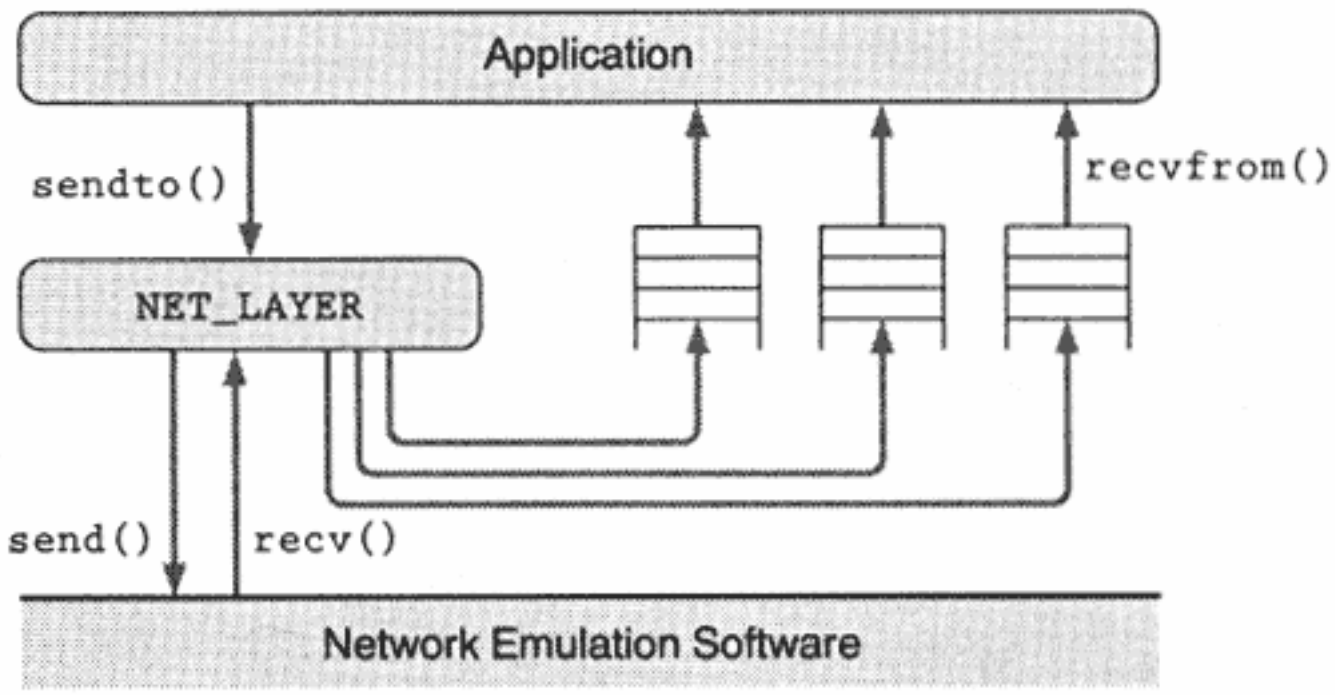

Figure 13.18
Message flow in the Commkit sockette implementation.

Fortunately, the second problem is solved by the first. By restricting the use of sockettes to the foreground process, the background queues BACKGROUND_2, BACKGROUND_3, and BACKGROUND_4 can be used to queue sockette messages (queue BACKGROUND_1 is used by the network layer). Whenever a successful bind() takes place, the resulting sockette number is used as an index into port_queue_lista list of the aforementioned background queues.

The arrival of a sockette message causes the network layer to perform two tasks:

1. Forward the incoming message to the queue associated with the sockette (one of BACKGROUND_2, BACKGROUND_3, and BACKGROUND_4). The message will be retrieved by the application when it reads from the sockette.

2. Send a 1-byte message to the foreground process (queue APPLICATION).

When the foreground process has completed processing and requires further information, it must call wait_event(). Wait_event() performs a recv() on the queue APPLICATION, thereby ensuring that the network layer can be invoked to handle any sockette messages. When control returns to the foreground process, either a keyboard character or the 1-byte network layer message is available. If the character is not from the keyboard, the foreground process can call recvfrom() to obtain the first pending sockette message.

13.5.5
Sockette Example

Three sockette function examples are found in sockex.c:

do_echo(). An example of echo server; any message that it receives is returned to the client (i.e., the sender of the message). Determining the client is a trivial operation, since the source of the message is supplied in the from address, obtained when recvfrom() is called. The from address is used as the destination address when the message is returned. The echo process binds to port ECHO_PORT.

Echo servers are often used when testing communications software since they have a well-known address and perform the basic server functions of receiving a request and issuing a response.

do_time(). A time server that returns the time of day to any client that requests it. As with the echo server, the client's address is available in the from address. Time() is called to get the time in seconds, ctime() is called to convert it to ASCII, and the time of day is returned to the client process. The time server process binds to port TIME_PORT.

Time servers are often found in network environments where machines do not maintain the time of day when they are shut off. The time server

allows these machines to initiate their time-of-day clock when they are first brought on-line.

`do_send_and_recv()`. Allows the user to send and receive sockette messages from any other process on the network (including itself). A series of commands allows the user to manipulate the sockette software. The commands the user can select are processed in `process_string()` and are as follows (all commands are terminated with a carriage return):

To bind to a port, the user types CTRL-B NUMBER, where NUMBER is any port number between 0 and 9. If the bind is successful, the sockette number associated with the port is displayed to the user; otherwise, an error message is issued. Neither the sockette number nor the hostport pair is kept, since the sockette number is supplied to the user for subsequent use.

To release a sockette (and hence its port), the user types CTRL-R NUMBER, where NUMBER is a sockette number between 0 and 2. If the sockette cannot be released, an error message is issued. The sockette number is obtained from the command CTRL-B.

Each sockette message must be sent with a source host-port pair. Since the user has three different sockettes from which to choose, the command CTRL-A NUMBER allows the user to select a specific sockette number. The NUMBER must be a valid sockette number (obtained using CTRL-B); otherwise, the selection will fail and an error message will be issued. The sockette number is assigned to `active_xmit` and is used in all subsequent `sendto()` s until the next CTRL-A.

The user can also specify the sockette from which messages are to be received. The command CTRL-W NUMBER causes the software to check for messages sent to the specified sockette number. If the process has not bound to the specified number, `recvfrom()` will fail continuously; however, a diagnostic is *not* generated.

If NUMBER is valid, `waiting_for_message` is assigned `TRUE` and `active_recv` is assigned the sockette number. `Active_recv` is used in all subsequent `recvfrom()s` until it is changed.

A sockette message requires a destination host-port pair before it can be transmitted. The user can specify the intended destination using CTRL-D HOSTNAME,PORT. If HOSTNAME and port are valid, all subsequent messages are sent to this destination. Error messages are issued if either HOSTNAME or PORT is invalid.

Once the remote sockette host-port pair is determined, it is stored in `rmt_sock`. All subsequent transmissions are to the socket indicated in `rmt_sock`.

Strings with their first byte other than the control characters specified above are stored in `rmt_sock` and transmitted to the destination (previously specified by CTRL-D), using the active sockette (specified by CTRL-A). If `sendto()` fails, an error message is generated.

13.5.6
Compiling and Running the Sockette Software

The sockette software functions with either of the local area networks and the wide area network. All foreground process software is in `sockex.c`, which requires `sockcode .h` to compile. The sockette and network layer software can be found in `sockette.c`. Both `sockcode.h` and `sockette.h` are needed to compile `sockette.c`. All of the sockette software is network-independent except for the code within the network layer which conditionally compiles code for whatever network is requested. The specific dependencies are given in `makefile`.

As an example, to make all three different networks, type the following:

```
C:\> make ringsox.exe
C:\> make bussox.exe
C:\> make wansox.exe
```

The resulting executables contain the network emulation software, the sockette software, and the sockette demonstration software. Note that the hardware configurations associated with each type of network must still be adhered to. To run the sockette software, type the name of the executable and the following parameters:

```
C:\> bussox LineSpeed Hostname Option {M or WS}
```

where:

`LineSpeed` is the line speed (50, 300, 600, 1200, 2400, 4800, or 9600).
`Hostname` is the name of the local host (as defined in the file `hosts`).
`Option` specifies the service that the software is to perform. One of

`E` or `e`. Invokes the echo service on this PC. Messages sent to `Hostname` with port `ECHO_PORT` are returned.

`T` or `t`. Invokes the time service on this PC. Time request messages sent to `Hostname` with port `TIME_PORT` result in the current time of day being returned.

Any other character causes the PC to enter send-and-receive mode.

`{M` or `WS}`if the underlying network is the ring network, `M` indicates that the station is the monitor station (there can be only one). If the underlying network is the WAN, `WS` is a value between 1 and 7 denoting the station's window size. The field is ignored by the bus network.

When execution begins, the screen is cleared. If either the echo service or the time service has been specified, the sockette address of the source process is displayed. Both of these services can be aborted by typing CTRL-C.

If the PC is in send-and-receive mode, the user can explicitly bind to one or more ports, send and receive messages, and release ports. The following

is an annotated example of a session between a user on host `marvin` (host address `12430`), the time server on `tardis` (host address `12345`), and the echo server on `zaphod` (host address `13001`). The startup commands for the various PCs are shown in Table 13.1.

The example is as follows:

First, the user must bind to a port, in this case, port 3 (note that there is no space between CTRL-B and 3):
User: CTRL-B3
bussox: `Sockette 0 bound to 12430 3`
The response from `bussox` indicates that all port 3 communications will take place through sockette number 0; no further references to the local host or port are required.

In this example, the user both sends and receives on the same sockette (since the time server and echo server return their responses to the source host-port indicated in the request). The test software must be initialized accordingly: the first step sets up sockette `0` for transmission; the second step sets up the same socket for reception:
User: CTRL-A0
bussox: `Sockette 0 is the transmission sockette`
User: CTRL-W0
bussox: `Sockette 0 is the reception sockette`

The user must also specify a destination; to obtain the current time of day from the time server bound to `tardis` port 7, the user first indicates the destination and then sends a message to the server. If the server and its host are active, the time should be returned:
User: CTRL-DTARDIS,7
bussox: `Remote selected: 12345 7`
User: TIME PLEASE!
bussox: `Message received: Tue Feb 13 14:24:37 1996`

Each time the user sends a message (i.e., any string without a command prefix), the sockette message is sent across the underlying network to the time server bound to `tardis` port 7. The time server displays the message

```
Time request from: 12430 3
```

Table 13.1 Sockette startup commands.

Host	Command
`marvin`	`bussex 1200 marvin x`
`tardis`	`bussox 1200 tardis T`
`zaphod`	`bussox 1200 zaphod E`

If the user sends several messages before the time server has responded, all messages are forwarded. As the time messages are received, they are displayed.

By changing the destination, the user can access the echo server, bound to port 2 on host `zaphod`:
User: CTRL-DZAPHOD,2
bussox: `Remote selected: 13001 2`
User: HELLO HANDSOME!
bussox: `Message received: Hello handsome!`
On `zaphod`, the echo server issues the message

```
Echo request from: 12430 3
```

If `zaphod` is not part of the network, the message will be sent but a reply will never be received.

Finally, if the active reception socket is not the same as the transmission socket, the responses from the two servers will not be displayed. Consider, for example, the effect of binding to port 7 on `marvin` and making the resulting sockette the reception sockette, when sockette 0 is still the transmission sockette:
User: CTRL-B7
bussox: `Sockette 1 bound to 12430 7`
User: CTRLW1
bussox: `Sockette 1 is the reception sockette`
User: PLEASE REPLY

The echo request is received by the echo server on `zaphod` and echoed, but the returned messages are never displayed. By changing the reception sockette back to sockette 0, this message is displayed:
User: CTRLW0
bussox: `Please reply`

The DELETE key allows the deletion of any invalid or unwanted characters.

13.6 Summary

This chapter considered some of the issues surrounding two closely related topics: the identification of entities other than hosts or nodes and the need for layering in communication systems.

Host (and node) identification is achieved by specifying a destination address in the network's frame. By itself, the destination host's unicast address cannot be used to uniquely identify an entity, such as a process, since a host may support tens or hundreds of processes. Instead, when transmitting a message to a remote entity, it is necessary to identify both the entity and its host. In some applications, such as electronic mail, the entity receiving the message may in turn distribute the message to other entities on the host.

Since it is unrealistic to expect a single process to handle functions such as mail distribution, the framing of packets, file transfer, and untold other tasks, many communication systems have been developed in a layered fashion. That is, the communication functions have been divided into a number of well-defined operations and assigned to a layer. Each layer offers the layer above it a number of services while using the services of the layer below it.

There is another demand for layering: hiding network- or manufacturer-specific functions, thereby permitting the interconnection of software developed on different machines and on different networks. For layering to work properly, it is necessary to develop standards that are agreed to by the people who use the layers. A number of well-known layered architectures have been proposed and implemented since the late 1970s, including the Open Systems Interconnection model, IEEE 802, and the DARPA protocol architecture.

Since the late 1980s, there has been explosive growth in the use of the Internet and the demand for client-server software. Although many applications seem novel, it is important to remember that ultimately the information to be accessed is simply a stream of bits that utilizes an underlying communication system. Understanding the concepts associated with layered architectures makes it easier to understand these applications.

13.7 Exercises

When testing the various networks using sockettes, remember that the rules governing the physical connections associated with each network must still be applied.

1. (A) Into which OSI layer should the following services be placed? (Justify each answer.)

    a. A time server.

    b. A service that performs byte swapping from Intel format to Motorola format (see Chapter 1).

    c. The X.25 data-qualifier bit (the Q-bit) (see Chapter 10).

    d. The functions of a ring network monitor.

2. (A) UNIX sockets reserve a number of ports (typically the first 512) for "well-known" services. Suggest reasons why this approach is taken.

When answering this exercise, consider the effect on a Commkit process if the process is arbitrarily bound to a reserved port number (such as the `ECHO_PORT`). What would happen if echo messages started arriving at the process?

3. (B) Set up a network (bus, ring, or wide area) consisting of two stations and the line analyzer. Run the echo process on one station and the send-receive process on the other. Transmit a message to the echo process. What appears on the line analyzer? Identify each field in the frame; remember, the host numbers are two bytes long while the port numbers are one byte long.

4. (B) Set up a network (bus, ring, or wide area) consisting of three stations (and the line analyzer, if possible). Run the echo process on one station and the `send_and_receive` process on the other two. Transmit a message to the echo process from each station. How does the echo process "know" which station is to receive the message?

5. (B) Set up a network (bus, ring, or wide area) consisting of two stations (one being the line analyzer). Run the `send_and_receive` process, binding to port 4, and then set up the sockette for both transmission and reception. Transmit a message to the `send_and_receive` process, port 4. Explain what happens. Does the message appear on the network? Explain.

6. (B) Set up a network (bus, ring, or wide area) consisting of two stations (and the line analyzer, if possible). Run the `send_and_receive` process on the other two stations and establish a sockette on both. Is it possible to communicate between the two stations? Compare this approach to message transmission using one of the test networks (such as `ringex` or `busex`). Which is more user-friendly?

7. (B) Set up a local area network (bus or ring) consisting of two stations and the line analyzer. Run the network test software (`ringex` or `busex`) on one station and the `send_and_receive` process on the other. Is it possible to transmit a message from the `send_and_receive` process to the station running the network test software? What appears on the test station when a frame finally does arrive?

Can a message be sent from the test station to the station running the `send_and_receive` process? Does the `send_and_receive` process receive the message?

8. (B) Compare the connection-oriented wide area network software (`wannet.c` and `wanex.c`) with the connectionless sockette software (`sockette.c` and `sockex.c`). What are the differences between the connection-oriented and connectionless software? How can `wanlow.c` support both connection-oriented and connectionless communications?

9. (C) If a process never reads its sockettes and messages keep arriving, the station eventually runs out of buffer space and crashes. Modify `sockette.c` to handle the situation in which a process never reads a socket. Put a threshold of 10 on the number of messages that will be held before the oldest message is discarded.

10. (C) A name server is a service that maps names into addresses, somewhat analogous to a telephone company's directory assistance. For example, a person can ask the operator for someone's telephone number; the operator attempts to find the number and gives it to the person. Implement a name server using sockettes. It should support three activities:

> A process can register its name and sockette address.
>
> Any process can query the server for the address of a given name. If the name exists, the address is returned; otherwise, an error indication is returned.

A process can remove its own name and address from the name server's list.

11. (C) In IP, the Internet-address-to-physical-address mapping table (`addr_to_physical` in sockettes) is updated dynamically by the Internet layer broadcasting an Address Resolution Protocol (ARP) packet whenever confronted with an unknown Internet address. The ARP packet contains the name of the unknown host; the host with this name is expected to respond with its IP address. Design an ARP protocol for sockettes and then modify `sockette.c` so that unknown host addresses are resolved and the `addr_to_physical` table is updated.

12. (C) Many programmers feel that message passing is an unnatural programming paradigm, preferring procedure calls. A programming paradigm known as remote procedure calls has been developed that permits programs on one machine to call a subroutine on another machine. The remote procedure call (RPC) is implemented as a local procedure call to a *stub* procedure that creates a message containing the parameters and sends it to a remote stub procedure. The remote stub unpacks the parameters and makes a local procedure call to the called (remote) procedure. Once the target procedure has completed its task, the results are returned by the remote stub in a message to the local stub. The local stub unpacks the message, updates the necessary parameters, and returns to the calling procedure.

Implement RPCs using Commkit `send()`s and `recv()`s. Issues that should be considered include how call by value is supported, how call by address is supported, and what happens if the remote procedure's host fails.

To test the software, write a remote procedure that accepts two numbers, adds them, and returns the answer in a third number.

14
Interconnecting Networks

14.1
Introduction

One of the biggest criticisms of the explosive spread and popularity of the PC is that the growth and choice of equipment is very often uncontrolled. For example, upper-level management within a company may decide that all employees will use specific machines, operating systems, and software packages. However, as employees learn more about the equipment, they may opt for different operating systems or software packages, or may even change machines entirely.

Although independent thought is admirable, it can lead to problems. For example, what happens if the payroll manager unilaterally decides to purchase an 802.3 bus network to interconnect all of the stations in the payroll office, while at the same time the personnel manager purchases an 802.5 ring network? Sooner or later, someone will want to copy information from a station in the payroll office to a station in the personnel officefor example, from station N to station D in Figure 14.1.

Since the stations are not on the same type of network, direct interconnection and subsequent information transfer is not possible because of electrical, packet-structure, address, and other protocol differences. Solutions such as physically transferring information on diskettes may be satisfactory on occasion, but if large volumes of information are to be transferred regularly, some other procedure is required.

Similar problems occur when two offices in different cities, each running its own local area network, need to communicate. If they are interconnected by a wide area network, the ideal solution is to transfer information across it. However, since the WAN probably does not support the protocol used by the LANs, nor does it achieve the speeds of the LANs, direct connection of the two is not possible.

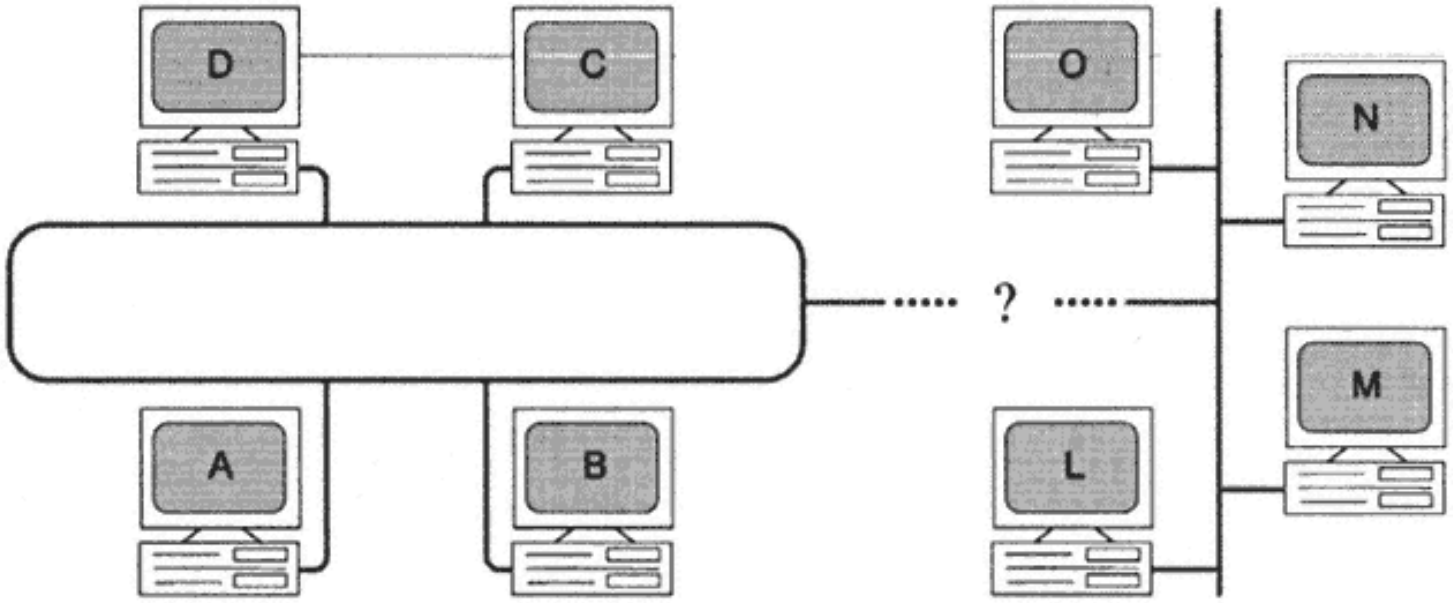

Figure 14.1
Interconnecting a ring to a bus.

Clearly, many of these problems will disappear if the networks are identical, since the packet structures and protocols are the same: simply interconnect the two (i.e., make a larger bus or ring). However, there are instances when it may not be physically possible or even advisable to allow all of the traffic on one network to propagate onto another. Consider the following situations:

Two lightly loaded 802.3 networks are to be interconnected. Combining them into a single, larger network may result in a heavy network load. The resulting network may offer less than satisfactory performance because of the additional network traffic.

Two local area networks that are more than one kilometer apart. The physical interconnection of the two may be impossible because of the maximum physical separation of stations allowed by the network.

A number of local area networks, each supporting several hundred stations. The interconnection may be physically possible, but the resulting number of stations may exceed the number allowed on a single network.

Two local area networks are used within an organization, one with sensitive product information, the other for use by the office staff. Combining the two and allowing the sensitive information on the general network may breach company security.

The above examples state the need for network interconnection, and at the same time outline some of the problems associated with it. If networks are to be interconnected, it is necessary to develop facilities that can convert between the different packet structures, addressing schemes, and other network idiosyncracies, as well as obey the protocols of the individual networks.

14.2
Background

Connecting a network to one or more other networks is possible as long as the device(s) supporting the interconnection can handle the various differences

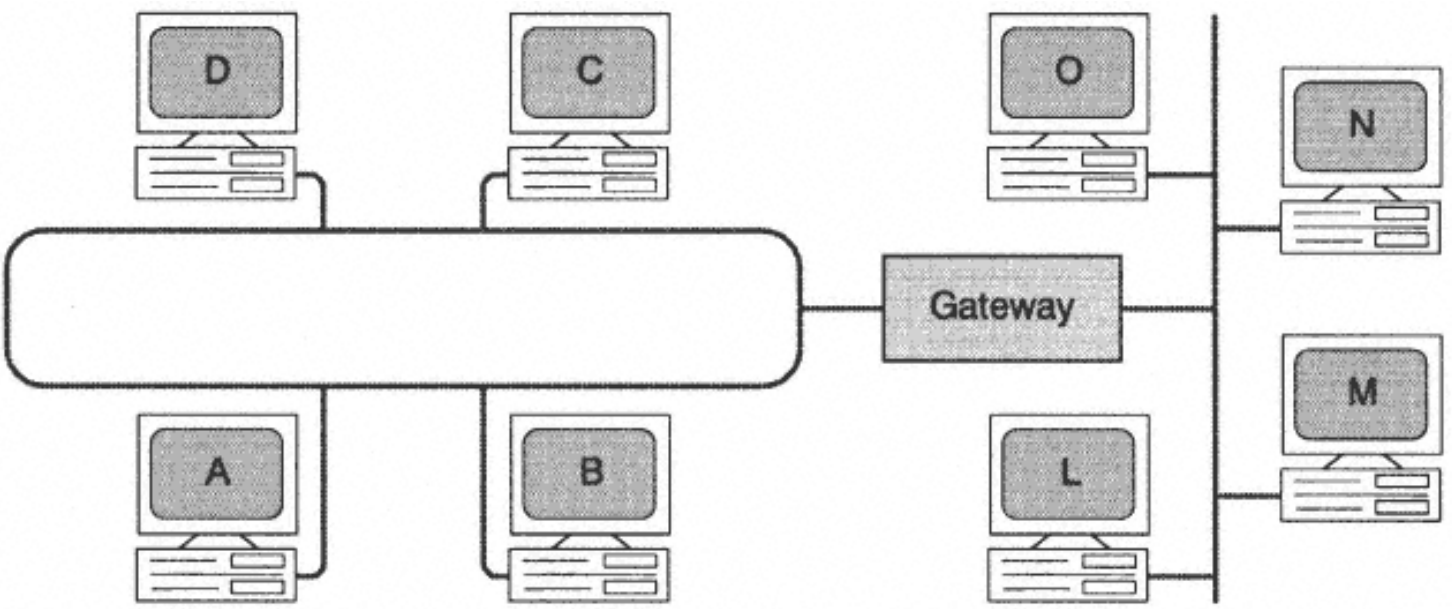

Figure 14.2
A gateway connecting a ring network to a bus network.

between networks. Although there are many types and makes of device to support network interconnection, manufacturers and users often employ the generic term gateway when referring to them. In Figure 14.2, a gateway connects a ring network to a bus network.

A gateway connects two or more networks while obeying the protocols associated with each one. For example, the gateway shown in Figure 14.2 must handle both the bus network and the ring network protocols. If the gateway is responsible for changing the packet structure on one network to conform to the rules of the other one, it is often referred to as a protocol converter.

In some situations, it may not be possible to have a single gateway (for example, if the networks are separated beyond the physical limits allowed by them, or if security dictates which stations are allowed on a network). As shown in Figure 14.3, it may be necessary to have a pair of half-gateways, each attached to a network and the other half-gateway.

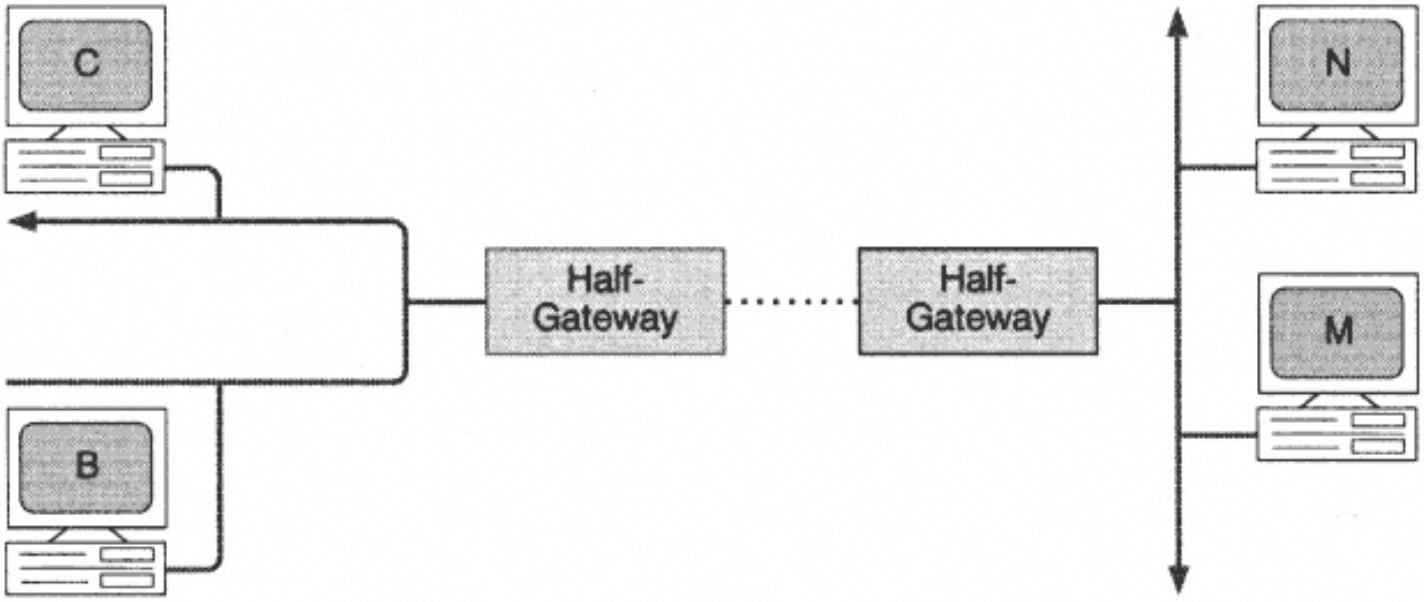

Figure 14.3
Two half-gateways interconnecting two networks.

Each half-gateway is connected to a single network and conforms to the standards associated with it. To send packets between the networks, the half-gateways must support a third, common protocol through which they communicate. Since administration of each half-gateway is left to the network to which it is connected, control over packets can be better maintained by the different organizations that run each network.

Other issues associated with network interconnection include

Fragmentation and reassembly. Maximum and minimum packet sizes can differ from network to network, potentially requiring the fragmentation or the division of the packet into smaller packets. The small packets are put back together in the larger packet using a process known as reassembly. When packets are fragmented, a mechanism such as sequence numbering allows the receiving entity to reassemble the packet in its correct order. The fragmentation can take place between the source and destination stations, using, for example, the smallest packet size of any of the intervening networks. When taking this approach, gateways are not required to perform any fragmentation or reassembly.

Alternatively, the gateways themselves can perform the fragmentation and reassembly. The destination station may still be expected to perform reassembly if the packet is too large for its network.

Connection-oriented vs. connectionless services. Gateways can potentially interconnect networks supporting connection-oriented (i.e., virtual circuit) servicestypically wide area networkswith networks supporting connectionless (i.e., datagram) networks. In these situations, it is necessary to decide upon the level and quality of service required by applications. If the communication is to be treated as a virtual circuit, the gateway may be required to support a reliable protocol atop the datagram service. However, an application using datagrams may find many of its packets being sent reliably across a virtual circuit, potentially adding to the cost of the overall communication.

Routing. Gateways can be required to perform routing if there are multiple paths through the network between stations. Routing algorithms such as those described for wide area networks in Chapter 10 are often used by gateways.

Other network differences. In addition to the above, the source network may support features that are not available on the destination network. A good example of this is synchronous mode on ring networks (see Chapter 12): a station transmitting a synchronous packet expects a reply from the destination within a specific period. If the destination station resides on a separate network, the delays associated with the intermediate gateway(s) forwarding the packet may make a reply within the allotted time period impossible to achieve.

If a gateway acknowledges the packet, this should not be interpreted by the source station as an indication that the destination station has

received the packet. In other words, higher-level protocols supporting end-to-end acknowledgments may be required to ensure that packets reach their destination.

As already mentioned, the term *gateway* is generic, applied to almost all network interconnection devices. However, in examining these devices on a layer-by-layer basis, one finds three types: repeaters (physical layer), bridges (data link layer), and gateways (network layer). Protocol conversion, another issue in internetwork communication, was discussed in previous chapters.

14.3
Repeaters

A repeater is a device that interconnects two homogeneous (i.e., identical) networks, making a single, larger network. Repeaters are intended to take the traffic from one part of the network and transmit it directly onto the other part without making any changes to the frame structure or the physical layer (i.e., OSI layer 1) protocol. A repeater, its relationship with two networks, and its associated layer are shown in Figure 14.4.

A repeater operates at the physical layer, amplifying or regenerating signals. Repeaters can help overcome segment length limitations and can be used to interconnect networks using different physical media (such as coaxial cable to fiber optic). A repeater connecting two bus segments passes the collisions between them. A ring repeater extends the size of the ring; the number of repeaters necessary in a ring equals the number of rings being connected. To avoid two or more stations receiving the same packet, all stations must be assigned unique addresses.

Finally, a repeater is not allowed to violate the physical limitations placed upon the network. Thus, two maximum-length Ethernets (2.5 kilometers) cannot be joined by a repeater to form a single 5-kilometer Ethernet, since this violates the Ethernet's cable length restrictions.

14.4
Bridges

A bridge is a device or layer of software that allows the interconnection of two or more local area networks at the MAC sublayer of the data link layer (OSI layer 2). The bridge structure is shown in Figure 14.5.

The resulting network is an individual LAN known as a segment. Each segment attaches to the bridge through a port, and a bridge can connect to many segments. Bridges act as store-and-forward devices, taking packets from

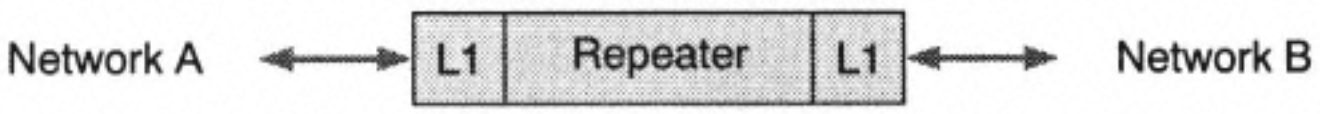

Figure 14.4
A repeater connects two networks at the physical layer.

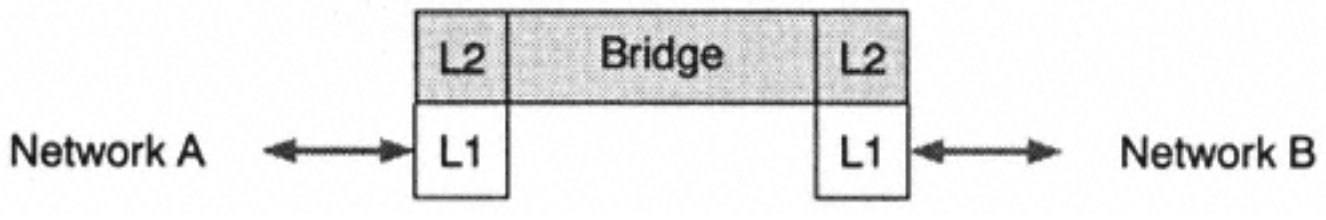

Figure 14.5
A bridge functions at the data link layer.

one segment and forwarding them onto others, following the rules associated with each segment.

Bridges offer a number of advantages over repeaters:

Different types of local area network can be interconnected. For example, a bridge can connect an 802.3 bus to an 802.5 ring. Packets destined for the ring must wait until the bridge is in possession of the token before transmission, while packets intended for the bus are transmitted when the bus is idle.

Distance limitations can be overcome. Two or more 802.3 bus networks can be connected by a bridge, and packets can be sent from one segment to the other through it. The bridge can give the illusion of a single, large network. But since the bridge operates at the data link layer, the physical layer requirements (such as distance limitations and electrical requirements) are not violated.

A single, large local area network can be divided into a number of smaller segments, all interconnected through one or more bridges. The bridge can help overcome the limitations associated with operating a single network that is approaching its physical limits (number of stations or maximum length).

The IEEE 802.1 standard defines two types of bridge: those that are transparent to the stations involved in the communication (transparent bridges), and those that require the stations involved in the communication to be aware of their existence (source routing bridges). Both types are now considered.

14.4.1 Transparent Bridges

To operate transparently, a bridge must function in promiscuous mode, receiving and inspecting all packets sent from each segment to which it is connected. Packets are addressed to the destination station rather than to a bridge. All stations must be assigned unique addresses.

In its simplest form, a transparent bridge forwards all packets received from one segment to all segments to which it is connected. This approach has a number of limitations, the most serious being flooding: all packets are sent to all segments, regardless of the destination. With all traffic allowed to propagate between the interconnected segments through the bridge, a series of lightly loaded segments can turn into a single, heavily loaded bridged network.

There may be a benefit in having the bridge perform some type of packet filtering. Ideally, the filtering allows only packets destined for a remote segment to pass through the bridge, thereby ensuring that the only traffic on a segment is either packets for stations on that segment or packets en route to a remote station being sent via a bridge.

In packet filtering, the bridge maintains a list of addresses for the stations associated with each segment to which it is connected; it can then selectively forward packets between segments by examining each packet's *destination* address. If the destination address is for a station on a segment other than the one from which the packet has been received, the packet can be forwarded to the destination's segment. To allow the station to make these routing decisions, a forwarding database is maintained, consisting of each known destination address and the segment from which the packet was received.

Using the example in Figure 14.6, if station S on segment 3 sends a packet to station E on segment 2, the bridge will forward the message to segment 2 because station E does not exist on segment 3 but it does exist on segment 2 (this is determined by the bridge examining the forwarding database).

Packets sent to stations on the *same* segment are not forwarded even though they are received by the bridge. For example, a packet from station D to station E will not be forwarded by the bridge because station E is on segment 2 (once again, the bridge determines this from the forwarding database).

If the bridge receives a packet with an unknown destination address, the packet is forwarded on all segments except the one on which it was received. For example, if station P sends a packet addressed to station X, the packet will be forwarded on segments 2 and 3, since address X does not exist in the forwarding database. If other bridges are attached to any of these segments, they will continue to forward the packet until it reaches station X. Having the

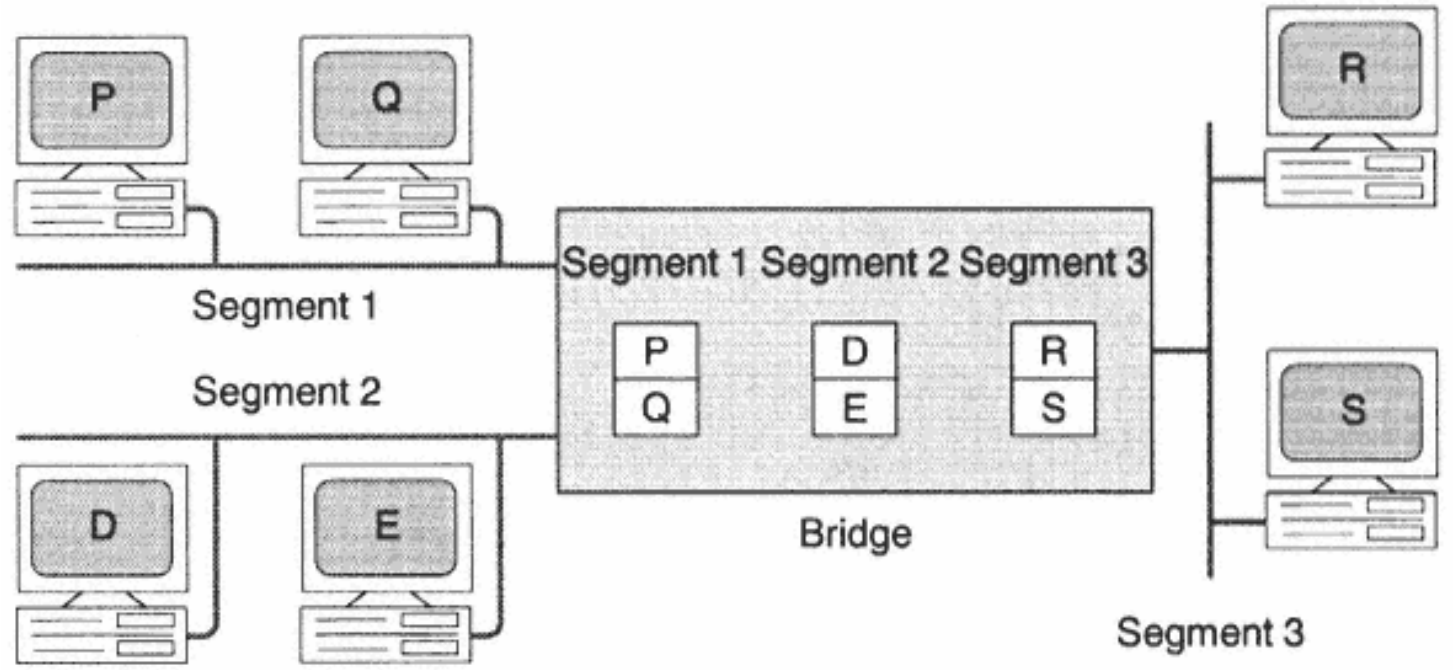

Figure 14.6
A transparent bridge.

bridge selectively forward packets ensures that the only packets on a segment (other than those en route to another bridge) are those intended for stations on that segment.

The forwarding database can be initialized in one of a number of ways:

The station addresses can be stored in a file that is read by the bridge when it is brought on-line. The problem with this approach is that the contents of the file are fixed, meaning that changing a station (i.e., removing it or having it change segments) requires that the file be explicitly altered, typically by the network manager.

Individual stations can explicitly inform the bridge of their existence. When a station joins or leaves a segment, the bridge can be informed, causing it to update its forwarding database. Although this eliminates the need for explicit database alteration by the network administrator, special protocols are required to allow the stations to inform the bridge of their status.

The bridge can determine the locations of the various stations on the segment by examining the *source* address of each packet. That is, whenever the bridge receives a packet from a segment, it immediately knows that there is a station on the segment with the address found in the packet. For example, in the configuration shown in Figure 14.7, packets sent on segment 1 will have the source address D or E, while packets sent on segment 2 will have the source address R or S.

By examining the source addresses associated with each packet, the bridge can construct the forwarding database associated with each segment; this is known as an intelligent bridge.

An intelligent bridge initially forwards *all* of the packets that it receives, "learning" the location of each station from the *source* address while building the forwarding database. Eventually, the bridge "knows" the location of each station on a segment and can filter the packets by examining the *destination* address.

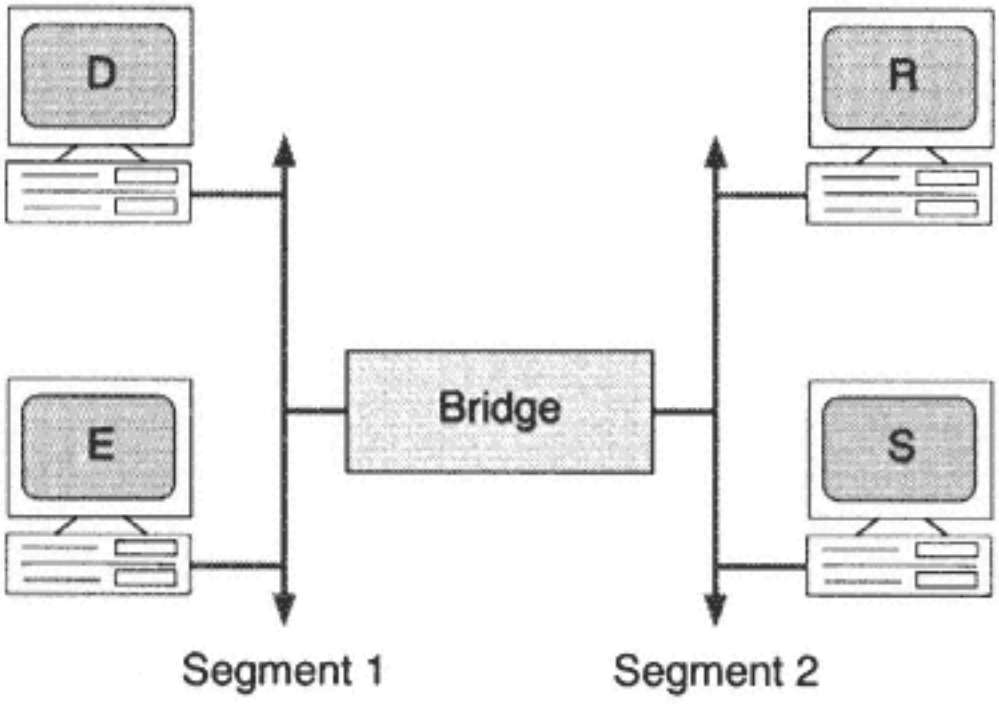

Figure 14.7
By examining the source of a packet,
the bridge knows a station's segment.

Should a station change segments or be removed from its segment, the bridge can also gain this information from the traffic. A station changing segments has its address on both segments; the bridge can then update its forwarding database on the basis of this information. In addition, the bridge determines when a station becomes inactive by periodically purging its forwarding database of stations that have not sent a packet within a fixed time (typically every few minutes).

IEEE 802 Transparent Bridge

The IEEE 802 transparent bridge is an intelligent bridge that performs packet filtering. Each 802 bridge maintains a forwarding database that consists of entries in the form *port*. the segment from which the packet was received; and *address*. the address of the station.

All packets received by the bridge are subject to two operations:

Bridge forwarding. When a packet is received, the bridge examines the forwarding database, comparing the list of station addresses to the destination address in the packet:

If the address does not exist in the database, the packet is forwarded on all segments to which the bridge is attached, except the one on which the packet was received.

If the address does exist in the database, the source port is considered. If the port from which the packet is received is the same as the port in the database, the packet is discarded; otherwise, it is forwarded to the segment specified.

Bridge learning. Once bridge forwarding is completed, the packet's address is compared to the list of addresses in the forwarding database. If the address is not there, both it and the source port are recorded in the database.

Inactive stations have their addresses removed from the forwarding database by the bridge. The amount of time an address is left in the database is set by the network administrator.

If reliability is an issue, some network administrators may institute a policy of having two or more bridges connecting a pair of segments. By having two paths through the network, an interesting problem arises: packets may be forwarded indefinitely. Consider the situation shown in Figure 14.8, in which segments A and B are connected by two bridges and station S on segment B attempts to send a message to station L.

When the packet addressed to station L is received by the bridges, both bridges attempt to forward the packet since (it is assumed) an entry for station L does not exist in either forwarding database (station L does not exist on either segment). Each bridge will then receive the packet forwarded on segment A from the other bridge and will forward it addressed to station L onto segment B. This cycle will continue indefinitely.

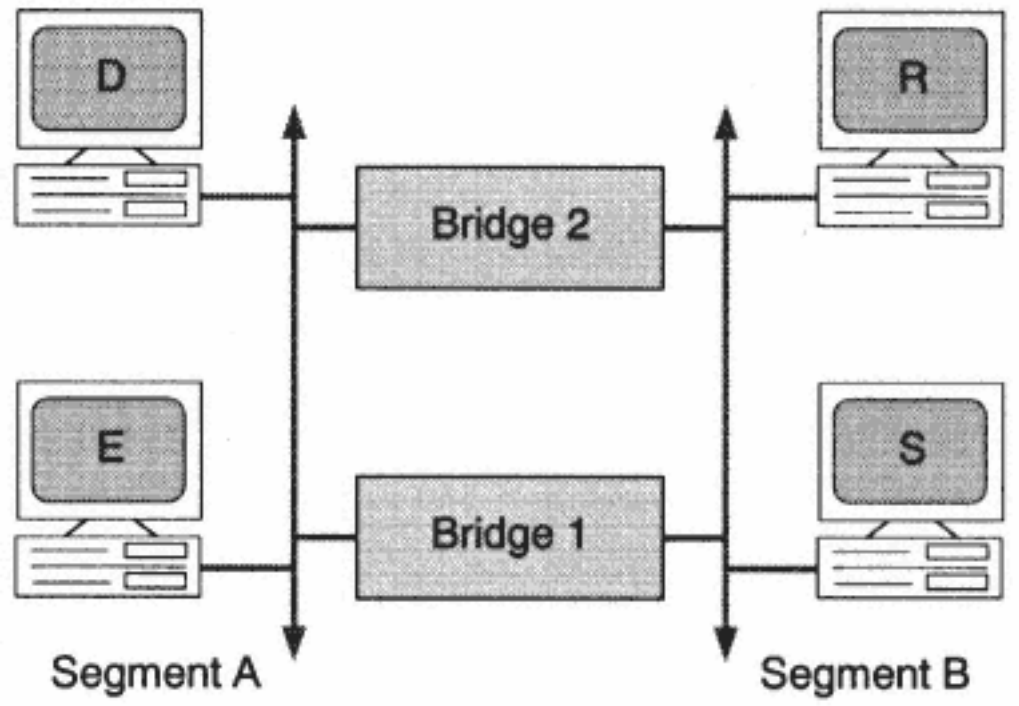

Figure 14.8
Interconnecting two networks with two bridges.

If the bridge learning and forwarding algorithms are to succeed, there can be only a single path between any of the segments making up the bridged network, regardless of the physical topology. Fortunately, it is possible to impose a spanning tree on a bridged network to ensure that there are no loops in the network. A spanning tree is a structure with one bridge as its root and other bridges restricting the flow of packets so that no loops exist within the overall network.

The 802 transparent bridge uses a spanning tree algorithm to generate spanning trees, which operate as follows. One bridge is chosen as the root of the spanning tree, typically the station with the lowest physical address. Then, all bridges determine the shortest path, from the segments to which they are connected, to the root bridge (i.e., the number of hops required by a packet to reach it). If two (or more) bridges connected to the same segment require the same number of hops, the bridge with the lower physical address is chosen, and the other bridge is *blocked*. (Blocking is a logical operation: the bridge still receives packets but simply discards those that are not to be forwarded.) The spanning tree algorithm functions dynamically, with each bridge periodically informing all other bridges of its location, neighboring bridges, and status. Should the status of a bridge or segment change (i.e., fail or become active), the spanning tree can be reconfigured.

For example, consider the LANs interconnected by bridges in Figure 14.9; since there are multiple paths through the network, packets can potentially cycle "forever."

If bridge B1 is taken as the root of the spanning tree, then the other bridges can determine the shortest path to the root bridge. The shortest path from segment 2 can be by bridges B2, B3, or B4. Assuming that B2 has the lowest physical address of the three bridges, B3 and B4 block (i.e., ignore) all packets sent to and from segment 2. The shortest (and only) path from segment 3 is via B4, which therefore will carry all packets destined for and

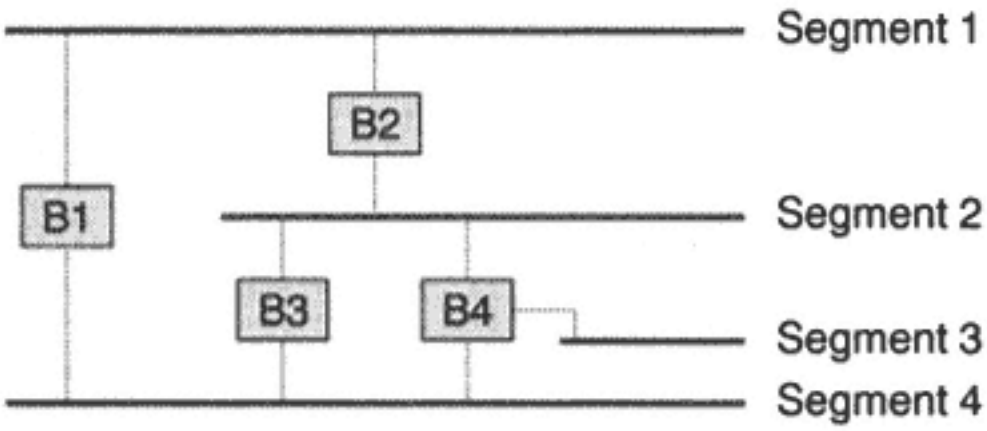

Figure 14.9
With multiple paths, a packet can circulate "forever."

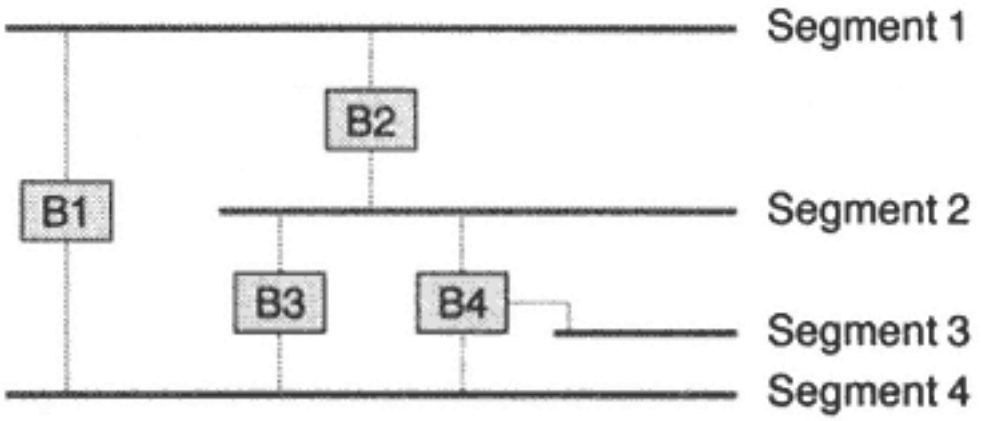

Figure 14.10
A spanning tree (from Figure 14.9).

received from segment 3. Bridge B3 blocks *all* packets and B4 blocks packets sent across segment 2. The resulting spanning tree is shown in Figure 14.10 (the dotted lines denote physical links that are currently blocked).

The bridged network is now free of all loops. Since the bridges are always in communication, a change in topology (for example, a bridge failure) will result in a new spanning tree.

14.4.2
Source Routing Bridges

Careful observation of the spanning tree in the previous section shows that the resulting bridged network, although free of possible loops, is not necessarily optimal. For example, a transmission from segment 3 to segment 2 results in the packet being sent via bridges B4, B1, and B2, even though the two segments are physically connected by bridge B4.

The designers of the IBM Token Ring recognized this limitation and developed an alternative to the spanning tree for bridged token rings known as source routing. In source routing, the source station, not the bridges, specifies the route of the packet. Not surprisingly, the route that the logical link control packet will take is stored within the frame by the source station (Figure 14.11).

Routing-Info is a variable-length field up to 18 bytes in length. The first two bytes, the routing control field, contain the following information:

Figure 14.11
The logical link control packet for source routing.

Whether the packet is a broadcast packet, intended for all rings.

The length of the Routing-Info field.

The direction of the packet, either from or to the source station.

An indication of the largest frame supported by the bridge for broadcast packets.

The remaining 16 bytes contain up to 8 route designators. The 16-bit routedesignator field consists of a unique 12-bit LAN number and a 4-bit bridge number identifying a specific bridge on a LAN.

Before a packet can be sent, the source must determine the route it will take. This is achieved by the source sending a discovery packet with the destination station's address to *all* token rings making up the network. As each discovery packet makes its way through the network, the route of the packet is recorded in the discovery packet by each bridge. Packets that reach the intended destination are returned by the destination to the source along the recorded route. The source can then choose which route to use for its communication from the discovery packets returned.

Once the route has been chosen, the source station sends packets with the high-order bit of the address set. This signals the bridge that the packet is intended for a remote token ring rather than the local token ring. The bridge inspects the Routing-Info field to determine the route.

Source routing has been adopted by the IEEE 802 committee for routing in interconnected 802.5 token rings. However, the source routing algorithm is not restricted to ring local area networks, but can be applied to other LANs as well as WANs.

14.5
Gateways

In many internetworking situations, it is not possible to use a repeater or a bridge because the networks to be connected support different physical properties, addressing schemes, and possibly even owners. A gateway (also known as a router) is a device operating at the network layer (i.e. OSI layer 3), that overcomes these limitations; the relationship between the gateway and the OSI layers is presented in Figure 14.12.

Since gateways are associated with wide area networks, they are, broadly speaking, of two types: those for connection-oriented services and those for connectionless services.

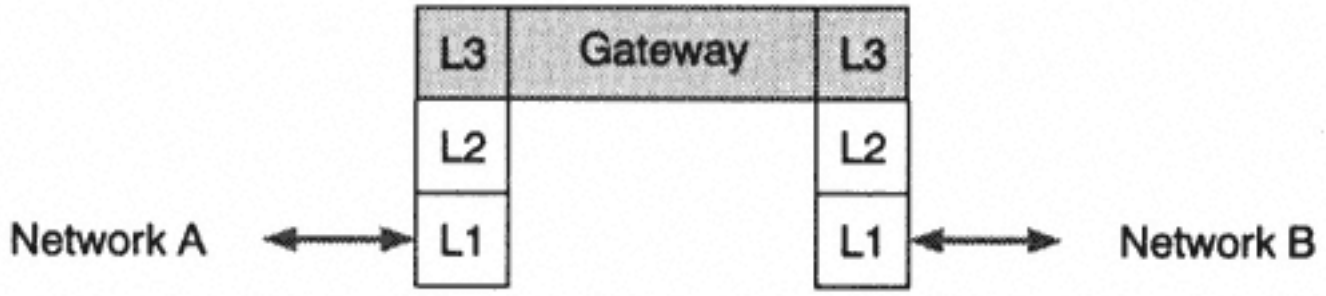

Figure 14.12
The relationship between a Gateway and OSI layers 1 through 3.

14.5.1
X. 75

The same arguments against the direct connection of dissimilar local area networks can also be applied to the direct connection of wide area networks, even if they share a common protocol such as X.25:

The internal protocols used by the network may differ. Connecting a node on one network to a node on another may not be possible because of physical connection (layer 1) differences or internodal protocol (layer 2) differences.

Since most public wide area networks are run by different organizations, both parties are usually reluctant to permit the other organization to connect one of "their nodes to our network."

Instead, most wide area networks that allow connection to other wide area networks communicate through a pair of half-gateways using a common gateway protocol. In the case of X.25, the gateway protocol is ITU X.75.

A connection-oriented protocol, X.75 is used when an application on an X.25 host (or PAD) specifies an international call in the call request packet. (The international call is indicated by an X.121 address specifying a network other than the current one.) The node to which the X.25 host is connected then creates a virtual circuit to the half-gateway node. The half-gateway creates another virtual circuit to the half-gateway on the remote wide area network. This one places a virtual circuit to the destination X.25 host (or another half-gateway, repeating the cycle). Once the virtual circuits are established across the various networks, communications may proceed.

The layer 3 protocols between a pair of X.25 hosts (denoted H) between nodes (denoted *N*) and a pair of half-gateways (denoted G) are illustrated in Figure 14.13.

The X.75 protocol is similar to the X.25 protocol. Additional features in X.75 are needed for routing and accounting.

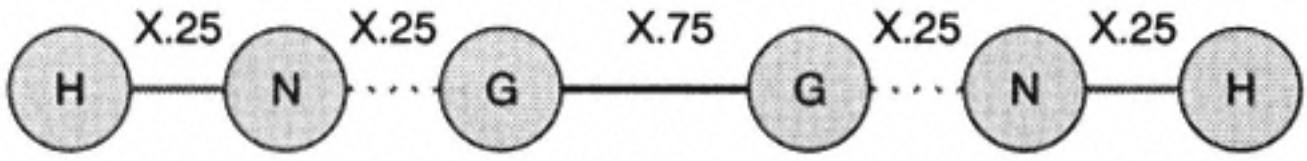

Figure 14.13
The relationship between X.25 and X.75.

14.5.2
IP Routing

In Chapter 13, it was shown how the Internet Protocol (IP) overcame different networks, addressing schemes, and protocols through the use of layering and by assigning each host a unique Internet address. When sending a message on the Internet, the application supplies a message, the destination IP address, and the port of the remote entity. The IP software takes the IP address, maps it into a unique physical host address, and then transmits the message, encapsulated in an IP packet, on the local network.

The Internet also supports the routing of messages between hosts on different networks; messages are sent to gateways that are responsible for the routing of packets between networks. Since each host is assigned a unique IP address, the gateway can determine the route of the packet by comparing the packet's destination address with the addresses in the gateway's routing table. Routing is independent of the high-level protocol in the packet (i.e., TCP/IP or UDP).

The Internet is huge, consisting of tens of thousands of sites and hundreds of thousands of hosts. If each packet were routed solely upon the value of the host IP address, the size of the routing table would be enormous and updating the routing tables would be a neverending task as hosts were added to and removed from the network. To overcome these difficulties, an Internet address is actually hierarchical, consisting of *two* parts: a network identifier and a host identifier.

When an application has a message to send to an entity on a remote host, the message and destination IP address are supplied to the IP layer as usual. The packet is encapsulated in a datagram, but since the network identifier indicates a network other than the local one, the datagram is sent to a gateway on that network. The destination IP address is the address of the destination network/host, *not* an intermediate gateway.

When the gateway receives the datagram, the network identifier is extracted from the packet's destination IP address and used to determine the datagram's route. If the network identifier indicates a network to which the gateway is directly connected, the datagram is forwarded to the destination on that network. However, if the gateway cannot reach the destination network directly, its routing table is accessed. Each entry in the routing table indicates the address of the *next* gateway on the path to the destination network. If the specified destination network is not in the routing table, the datagram is forwarded to a default gateway. This cycle is repeated until the datagram reaches a gateway connected to the destination network. Note that the source and destination IP addresses are never altered by the gateways; the IP addresses are used for routing purposes.

The Internet supports *three* classes of Internet address, all in the form `net-id, host-id`. The number of networks associated with each class is shown in Table 14.1. Figure 14.14 shows how a gateway can distinguish be-

Table 14.1 Internet address classes

Class	Number of Hosts
A	More than 65,535
B	256 to 65,535
C	Less than 256

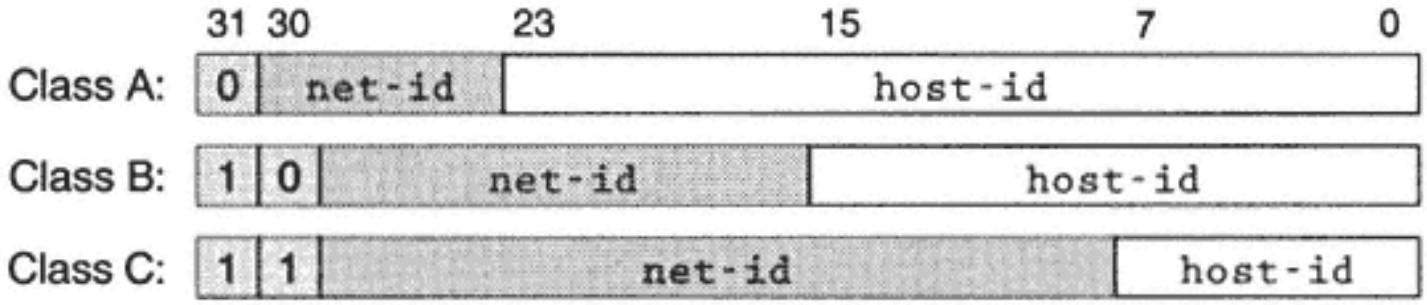

Figure 14.14
IP address structures.

tween the different address structures by examining the high-order bits of the IP address classes.

IP addresses are normally written in dot notation: each byte making up the 32-bit address is separated by a period, in the format `m.n.o.p`. Using the IP dot notation for address representation, class A networks have values `1.0.0.0` through `127.0.0.0` (the ARPANET is 10.0.0.0), while class B networks are in the range `128.0.0.0` through `191.255.0.0`; and class C networks are `192.0.0.0` through `255.255.255.0`.

Each IP gateway is responsible for maintaining its own routing table which can be initialized from static storage, or the gateway can communicate with other gateways, dynamically building the routing table. Only a subset of gateways maintains a list of routes to all networks; these gateways are known as core gateways and connect directly to the ARPANET. The remaining gateways are known as noncore. Routing tables are updated by the gateways exchanging routing information on the distances between gateways in terms of hop-counts.

14.6 The Commkit Transparent Bridge

Commkit has been supplied with software that emulates a transparent bridge, allowing packets to be sent between two local area networks (e.g., the bus or the ring) using the PC's parallel port. For example, in Figure 14.15 a pair of 2-station Commkit networks (one a ring and the other a bus) are connected by the parallel ports of two of the PCs.

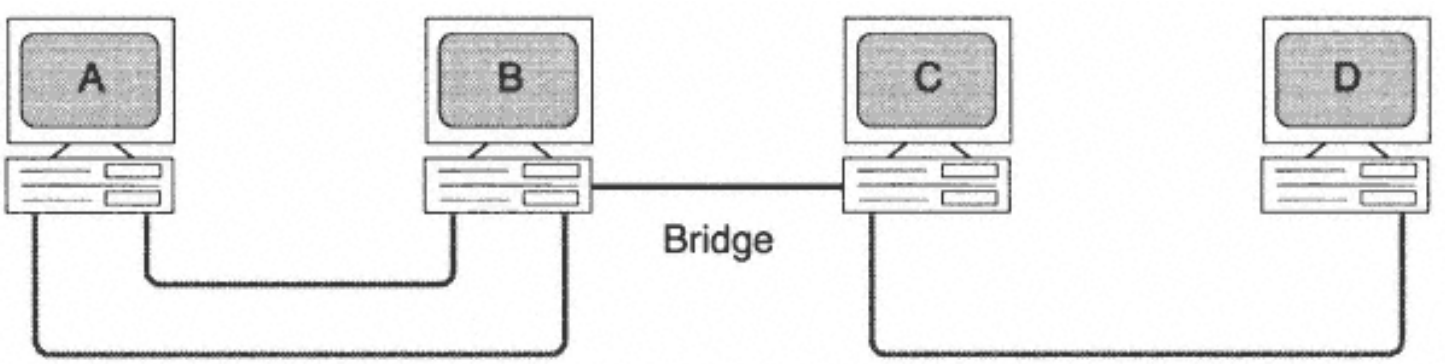

Figure 14.15
Interconnecting Commkit networks with the Commkit bridge.

In the figure, PCs B and C are connected to the ring and bus network, respectively, as well as directly connected by a cable through their parallel ports. The bridge is actually a pair of half-bridges, which are identical except for the underlying LAN software. A pair of background processes control the routing and formatting of packets. The foreground application allows the user to query the state of each half-bridge.

14.6.1
Hardware Requirements

The Commkit bridge software is written for unidirectional parallel ports, the type supplied with the IBM PC for its printer adapter. Normally, if a printer is connected to the PC, data is sent to it through the data register, port `0x3BC`, and the printer status is returned through the status register, port `0x3BD`. The control register is port `0x3BE`.

Simply connecting the printer ports, using a straight-through 25-pin RS-232-C cable, and writing data to the data register on either PC will not result in the transfer of data, since the data registers support output only and cannot be read from. Instead, the cable must be rewired so that the output from the data register on one PC appears as input on the other PC's *status* register. The pins to be exchanged are 2 and 15; 3 and 13; 4 and 12; 5 and 10; and 6 and 11. The swapping should be done on only one connector. Electrically, this should not be a problem since the data register pins produce 2.6 mA, while the status register pins can handle up to 7 mA.

The results of these changes are that data sent to the data register will arrive on the status register (note that the parallel port does not supply data on the lower 3 bits of the status register; these bits are always set), as illustrated in Figure 14.16.

Data received on pin 11 (bit 7 of the status register) is always the inverted value of what was sent (i.e., a 0 is received as a 1 and a 1 is received as a 0). Thus, if a cable is wired according to the above instructions and plugged into the parallel ports on a pair of PCs, then writing `0x1F` to the parallel port data register on one PC will result in `0x7F` being read from the status register of the other PC.

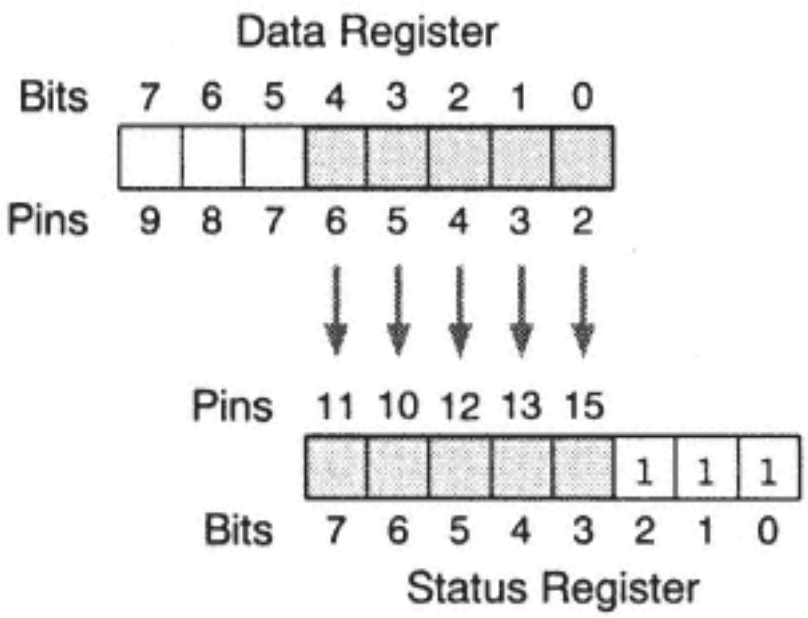

Figure 14.16
Parallel port wiring for the Commkit bridge software.

If parallel port interrupts are to be supported, bit 7 of the 8259 Interrupt Controller mask must be cleared, the parallel port interrupt vector (number `0x0F`) must contain the address of the interrupt handler, and the parallel port control register must be set to `0x10` to enable interrupts. A parallel port interrupt occurs when there is a high-to-low transition on pin 10 (bit 6) of the status register (this is achieved by asserting a low-to-high transition on bit 3; the parallel port inverts the signal).

The parallel port interrupt is *not* the same as the UART's. In the UART, an interrupt will occur after the transmission or reception of a character; whereas in the parallel port, an interrupt will occur only when pin 10 of the status register experiences a high-to-low transition.

14.6.2 The Bridge Software

The bridge software consists of two parts: the parallel port software for the transfer of individual bytes across the cable, `parlport.c`; and the software needed to control the flow of information between the bridges, `bridge.c`.

The Parallel Port Protocol

The parallel port protocol software, `parlport.c`, is responsible for the transfer of bytes between two PCs using the parallel port. A quick examination of the connections between the data and status registers reveals that, at most, four bits are available for a data transfer (bits 0, 1, 2, and 4), since the status port will accept only five input bits, one of which must be used to signal an interrupt (bit 3).

As there are four bits that can be used for data transfer, a byte can be sent from one PC to the other as a pair of nibbles. A total of five bits are sent: four data bits, the nibble, and the fifth bit (bit 3) to signal an interrupt. By sending each 5-bit sequence *twice*. the first time with bit 3 cleared and the second time with bit 3 set, an interrupt will occur on the receiving PC.

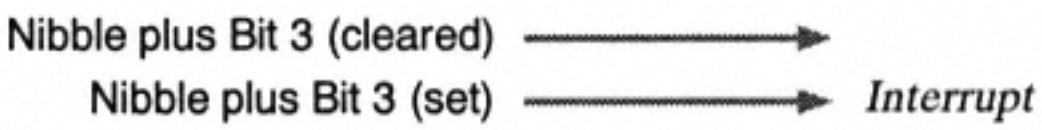

Figure 14.17
Interrupt signaling using a pair of nibbles.

Figure 14.17 shows how the receiving PC can read the status register to obtain the value of the five bits (bit 3 will be set).

The above operation must be performed for every nibble sent if an interrupt is to occur on the receiving PC. To send a byte, the cycle must be repeated for both nibbles making up the byte.

`Pp_out(ch)` sends the lower five bits of `ch` to the remote PC. `Ch` is written to port `PP_DATA` twice, the first time with bit 3 cleared, the second time with it set, thereby causing an interrupt on the remote PC.

This raises an interesting problem: how fast should the byte be sent? If the pair of nibbles making up the byte are sent too quickly, the receiving PC may not have sufficient time to process them; however, if they're sent too slowly, throughput may suffer. Since an interrupt occurs each time a nibble is sent, the PC receiving the nibble can respond with its own nibble (this and subsequent references to nibbles assume that when a nibble is sent, it is sent as a 5-bit sequence, causing an interrupt on the remote PC). In Figure 14.18, PC 1 is to send a byte to PC 2.

If full-duplex transfer is to be permitted (that is, both PCs can transmit bytes simultaneously), it is necessary to distinguish between nibbles that make up bytes and nibbles that are for acknowledgments, by sending a special nibble, analogous to the start bit, prior to the transfer of the byte. In the example shown in Figure 14.19, PC 1 is to send a byte to PC 2.

The byte transmission and reception algorithms are implemented as a pair of finite state machines in `parlport.c`, one for transmission (`pp_xmit()`) and the other for reception (`pp_recv()`).

The transmission state machine consists of four states, indicated by the value of `pp_info.xmit_state`, and it is implemented in `pp_xmit()`. `Pp_xmit()` is called after `pp_recv()` has processed the incoming interrupt. `Pp_recv()` supplies a single parameter, `reply_needed`, which indicates whether a reply to the remote PC in the form of an acknowledgment is required. The transmission states are as follows:

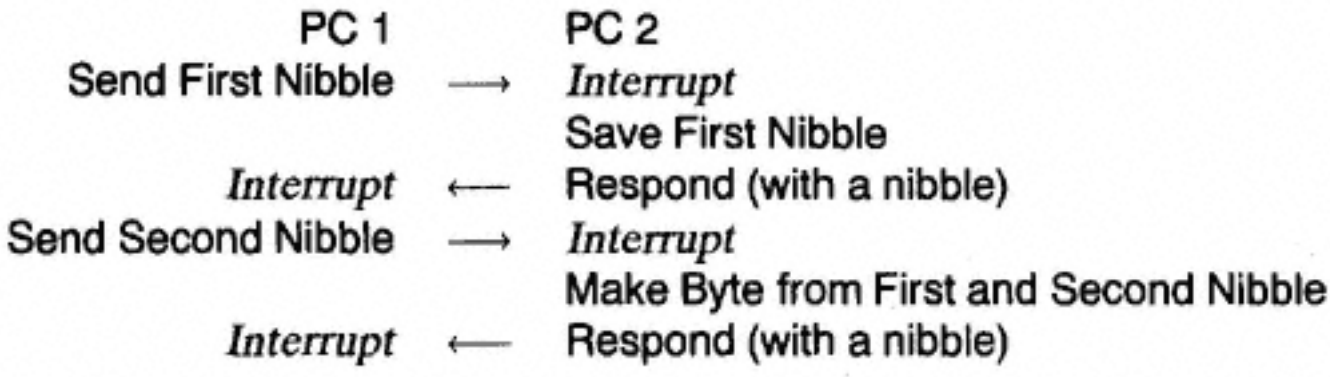

Figure 14.18
Byte transfer using pairs of nibbles.

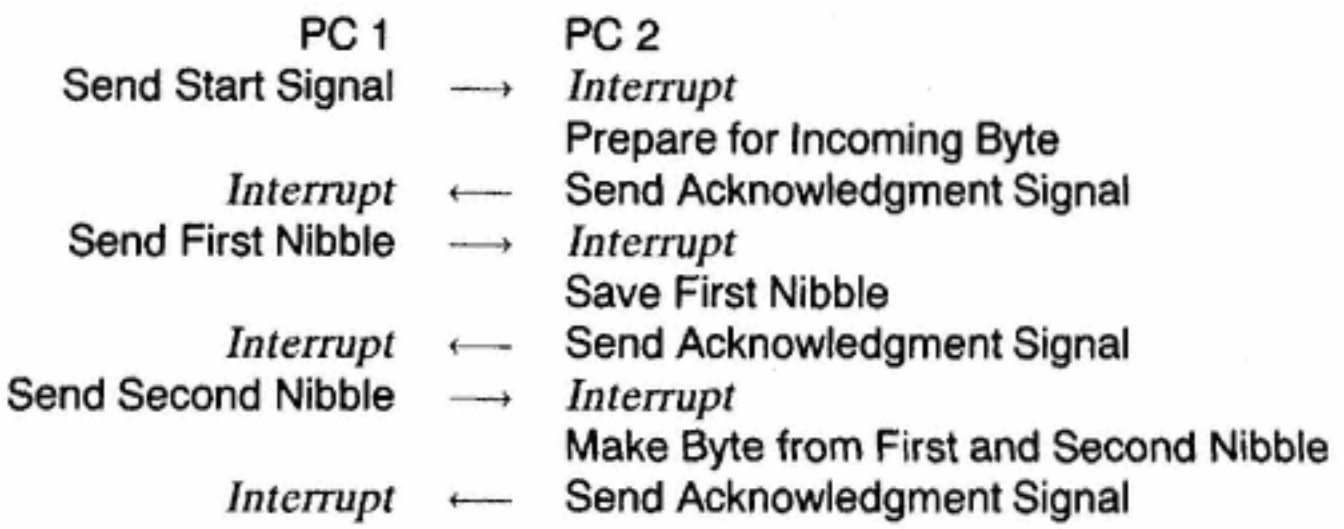

Figure 14.19
Distinguishing between acknowledgment nibbles and data nibbles.

`SEND_IDLE`. The initial transmission state, returned to after a byte has been sent to the remote PC. If there is a byte to be transmitted, `pp_info.data_to_go` is `TRUE`, causing a `START_SIGNAL` to be sent by a call to `pp_out()`. The byte to be sent is pointed to by `pp_info.x_ptr`. The state changes to `SEND_NIB1`.

If `pp_info.data_to_go` is `FALSE`, an acknowledgment signal can be transmitted if requested by `pp_recv()`. The state does not change.

`SEND_NIB1`. Sends the low-order nibble of the byte pointed to by `pp_info.x_ptr`. Figure 14.20 shows how the 4-bit nibble is encoded into five bits. Instead of shifting and masking the various bits, an array, `out_nibs`, is used, which is indexed by the nibble, producing the corresponding bit pattern that would be produced by shifting and masking. In all cases, the resulting bit pattern has bit 3 set. For example, a nibble value of `0x00` produces `0x08`, while a value of `0x0E` produces `0x1E`. The encoded nibbles are transmitted using `pp_out()`.

Once the encoded nibble is sent, the state changes to `SEND_NIB2`.

`SEND_NIB2`. Identical to `SEND_NIB1` with the following exceptions: the high-order nibble is encoded for transmission (using `out_nibs`), and the next state is `WAIT_ACK`.

`WAIT_ACK`. Entered after the second nibble is sent and the last acknowledgment is received. The receive state changes to `SEND_IDLE`, and `pp_info.xmit_done` is set `TRUE` to indicate that the transmission has completed.

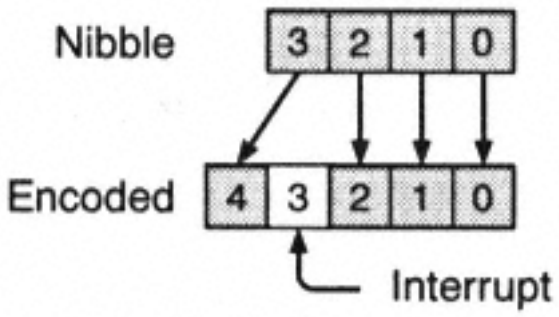

Figure 14.20
The encoding of the 4-bit
nibble into five bits.

In each of the above states that handle the transmission of a byte (i.e., all states except WAIT_ACK), a timer is started by calling start_pp_clock(). If a response from the remote PC is not received within two clock interrupts, pp_xmit() is called as if an acknowledgment had been received, causing the transmission to continue. If the resulting data value is invalid, higher-level software must resolve the error. The clock is turned *off* in state WAIT_ACK.

A transmission is never started unless the receive state machine is in state WAIT_BYTE. This policy was instituted after a number of tests showed that data could be lost if data transmissions were started while acknowledgments were being returned.

Pp_recv() is called by pp_protocol() (see below) whenever a parallel port interrupt occurs; that is, data is received by the parallel port status register. Pp_recv() shifts the encoded 5-bit status register value right by three, giving it a value between 0 and 31; if no errors have occurred on the channel, this value is the same as the original encoded value. How the encoded value is interpreted depends upon the state of the receive state machine. The receive state machine consists of three states (indicated by the value of pp_info.recv_state):

WAIT_BYTE. The initial receive state, returned to after a byte is constructed from a pair of nibbles. If the shifted status register has a value of START_SIGNAL, the receive state changes to WAIT_NIB1. Any other value is assumed to be an acknowledgment signal (ACK_SIGNAL) and does not cause a state change.

WAIT_NIB1. The four data bits of the register are stored in the low-order nibble of the byte pp_info.r_data. The value of the nibble is obtained by treating the value of the encoded five bits as a subscript into the array in_nibs. The state changes to WAIT_NIB2.

WAIT_NIB2. The four data bits of the register are stored in the high-order nibble of byte pp_info.r_data. Once again, the value of the nibble is determined from the array in_nibs. The state returns to WAIT_BYTE.

In all of the above states, once the encoded register value is processed, pp_xmit() is called with a Boolean value indicating whether the transmission state machine is to respond with an acknowledgment (ACK_SIGNAL), a start signal (START_SIGNAL), a data nibble, or not at all, depending upon the receive state and the transmission state. If a byte is received, a response is required to ensure that the remote PC can continue to transmit. However, if a byte is not received, responses may still be required if the local PC is transmitting. Pp_info.ack_count indicates the number of acknowledgments that are expected. If limits are not placed on the number of acknowledgments, a continuous stream will be exchanged by the PCs, seriously degrading the performance of each one.

The remaining three parallel port functions are as follows:

`pp_init()`. Initializes the parallel port software. If the parallel port software is being used in an application, `pp_init()` should be declared as an external and called during the initialization phase.

`do_pp_protocol()`. Called by `pp_protocol()` after each parallel port interrupt, since it is possible that a byte has been received or transmitted. These conditions are indicated by Booleans `pp_info.recv_data` and `pp_info.xmit_done`, respectively. If a byte has been received, the value in `pp_info.r_data` is forwarded to the process associated with queue `BACKGROUND_2`.

When a byte has been successfully transmitted, a check is made to determine whether there is more to transmit (indicated by `pp_info.x_count` having a non-zero value). Should there be nothing more to transmit, `recv()` is called and the queue `PPIH` is checked. If the queue has a message that is available for transmission, `pp_info.x_count, pp_info.x_buffer`, and `pp_info.x_ptr` are updated; otherwise, control returns to `pp_protocol()`.

In all cases, if there is a byte to be sent (pointed to by `pp_info.x_ptr`), `pp_info.data_to_go` is set to `TRUE`. If the receive state machine is in the `WAIT_BYTE` state, the transmission is initiated by a call to `pp_xmit()`.

`pp_protocol (code, ch)`. Should be called, if parallel port interrupts are supported, whenever a parallel port interrupt occurs. Parallel port interrupts are signaled by a device value of `PPIH` in `low_level()`. The code (one of `RECVDONE` or `MSG_AVAIL`) and the associated data can be supplied to `pp_protocol()` by `low_level()`.

The value of the status register is supplied in `ch` whenever a code value of `RECVDONE` is received. `Ch` is then supplied to `pp_recv()` for processing; once completed, `dopp_protocol()` is called.

When a message is available for transmission (signaled by a `code` of `MSG_AVAIL`), a check is made to determine whether the transmission can begin immediately by checking `pp_info.x_count` and `pp_info.recv_state. Pp_xmit()` is called to start the transmission. Clock interrupt signals (indicated by a `code` of `CLKIH`) are also required when the transmission state machine does not receive an interrupt. Accordingly, clock interrupts must be forwarded from `low_level()`.

The Bridge Software

Commkit supports a transparent bridge that connects two local area networks. The bridge is implemented as a pair of half-bridges, each of which is constructed as in Figure 14.21, regardless of the underlying local area network:

Network. Either of the LAN emulators, `buslow.c` or `ringlow.c`. The network software is identical to the software described in the section on local

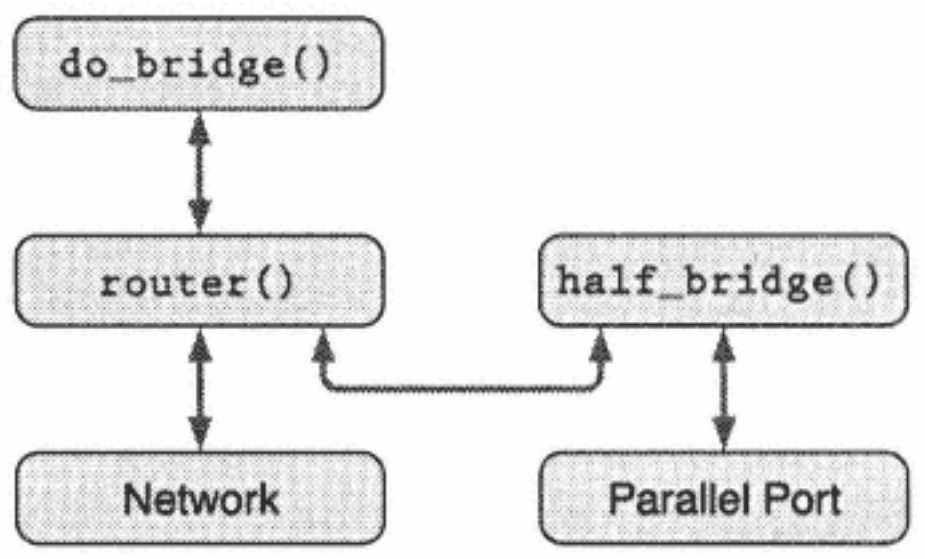

Figure 14.21
Message flow in the Commkit bridge software.

area networks, with two exceptions: first, the variable `promiscuous` is set `TRUE` by the bridge software, ensuring that all network frames are accepted by the station; second, the token `BRIDGE` is defined, so that parallel port interrupts are routed to the parallel port software.

Parallel Port. Described in the previous section. All messages from the `half_bridge()` are transmitted to the remote PC, while each byte received from the remote PC is forwarded to the `half_bridge()`.

`router()`. Packets received from either the local network (i.e., the network to which the half-bridge is attached) or the remote network are subject to bridge forwarding and bridge learning by `router()`, found in `bridge.c`. The incoming packet is taken from queue `BACKGROUND_1` and stored in `net_msg`.

Bridge forwarding involves extracting the destination station identifier from the packet and determining the destination network. Since both local area networks use the same packet structure, access to the packet addresses is network independent (`net_msg[0]` for the destination address and `net_msg [1]` for the source address). The destination network is obtained from the array `fwding_db`, using the destination address as a subscript (`fwding_db` is an array of 256 elementssufficient for all possible station addresses). An element in `fwding_db` has one of three values: the queue number of the local network emulator (`NETWORK`), the queue number of `half_bridge () (BACKGROUND_2)`, or 1, indicating that the network for this address is as yet unknown.

Since there are only two possible networks, routing decisions are as shown in Table 14.2.

Bridge learning consists of determining the source of the message and updating `fwding_db` with the queue number of the process that sent the packet (either `NETWORK` or `BACKGROUND_2`).

`half_bridge()` . Packets destined for the "other" network are forwarded to the `half_bridge` from either the remote `half_bridge()` (via the parallel port software) or the local `router()` (`half bridge()` is also found in `bridge.c`). Packets sent across the parallel port are enclosed in

Table 14.2 Routing decisions for a 2-network bridge.

Packet Source	Destination Network (from `fwding_db`)		
	`NETWORK`	`BACKGROUND_2`	Unknown
`NETWORK`	Discard	Send to `Background_2`	Send to `BACKGROUND_2`
`BACKGROUND_2`	Send to `NETWORK`	Discard	Send to `NETWORK`

a frame delimited by the bytes *STX* and *ETX; DLE* insertion is required to avoid confusing a data *ETX* with the end-of-frame *ETX*.

Frames from the parallel port are received one byte at a time by `half_bridge()`. These bytes are assembled into the array `remote_msg` and forwarded to `router()` when the frame-delimiting *ETX* is found. All framing bytes and inserted *DLE*s are discarded. If the trailing *ETX* is lost and the number of bytes is about to exceed the limit allowed by `remote_msg`, the reception is aborted.

Packets from `router()` are received in their entirety and must be enclosed in a frame prior to transmission. Since *DLE* insertion is also necessary, the packet from `router()`, `bridge_msg`, is copied one byte at a time into the array `pp_msg`. Once the copying is completed, the packet is sent to the parallel port for transmission.

`do_bridge(). Do_bridge()` (found in `bridge.c`), is the foreground process associated with queue `APPLICATION` and is intended to allow the user to monitor the state of the local half-bridge. The following commands are recognized by `do_bridge()`:

`R` or `r`. Displays the number of bytes sent to each of the active stations by the half-bridge. Each station's total is stored as an element of array `r_bytes`.

`S` or `s`. Displays the port number associated with each active station; the port numbers are the queue numbers of the different networks or the queue number of the half-bridge (`BACKGROUND_2`). The values are obtained from the forwarding database array `fwding_db`.

`T` or `t`. Displays the number of bytes sent by each of the active stations and forwarded to the remote network. The totals are stored in the array `x_bytes`.

14.6.3 Compiling the Bridge Software

The Commkit bridge is made by linking the following modules: `bridge.obj, parlport.obj, commkit.obj, srcalls.obj`, and `crc.obj`, with one of the local area network modules `buslo.obj` or `ringlo.obj. Buslo.obj` (and `ringlo.obj`) are versions of `buslow.c` (and `ringlow.c`) that have been compiled with the token `BRIDGE` defined. The two bridges that can be

made from these object files are defined in the `makefile` as `busb.exe` and `ringb.exe`.

For example, to make both bridges, (`ringb.exe` and `busb.exe`), type the following:

```
C:\> make ringb.exe
C:\> make busb.exe
```

14.6.4
Running the Bridge Software

Both bus and ring bridges are available in executable form on the Commkit diskette; they can also be created using the `make` utility (see above).

Since the bridge software is independent of any underlying network, it is possible to connect different network types. For example, the bridged network shown in Figure 14.22 can be built out of Commkit's ring, bus, and bridge software (the cabling rules discussed in previous chapters and the special parallel port cable connections must still be observed). The stations will run the software shown in Table 14.3.

Neither `busb.exe` nor `ringb.exe` supports transmission on its respective network, since both permit the querying of various bridge statistics: `R` displays the number of bytes sent to each active station; `S` displays each active station and its network; and `T` displays the number of bytes received from each active station. CTRL-C will abort the execution of a bridge.

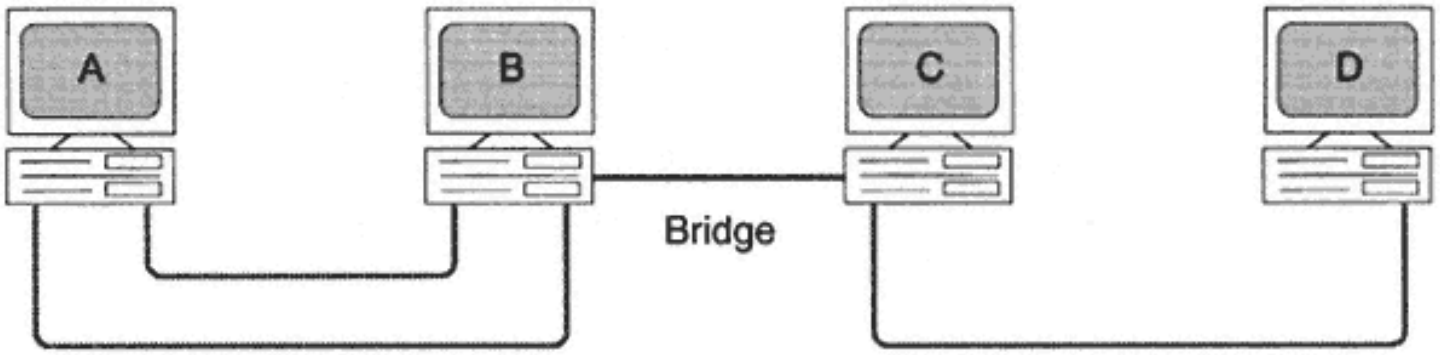

Figure 14.22
Interconnecting a Commkit ring to a Commkit bus.

Table 14.3 The software necessary for Figure 14.22.

PC	Software
A	`ringex 1200 A`
B	`ringb 1200 B M`
C	`busb 2400c`
D	`busex 2400 D`

It is advisable to run `ringb.exe` as the ring monitor, since this allows one of the remaining ring network stations to act as a line analyzer. For example, in the configuration in Figure 14.22, PC A can be run as a line analyzer, showing the traffic on the ring. This permits verification of packet forwarding across the bridge. A station on the bus network can also be run as a line analyzer.

Messages can be sent between stations on the same network or between stations on different networks. For example, a user on station D could send a message to the user on station A by typing

```
AD Nod your head if you receive this message!
```

Once a message has been sent, the bridges can be queried with respect to the active stations or the number of transmissions. Packet filtering takes place as soon as a station's identifier is entered into the forwarding database. For example, if station D sends a message to station A and then sends a message to itself, the message is not forwarded since the bus bridge "knows" that station D is on the bus network.

The addresses assigned to the bridges are not used by them, since bridges handle only the routing of packets and operate in promiscuous mode.

14.7 Summary

Ideally, all the services required by an application will reside on the application's network. However, because of equipment costs, physical distances, and organizational differences, the service required by the application may reside on a different network. In order for the application to access the service on the remote network, there must be facilities to allow the interconnection of the different networks. Broadly speaking, there are three such approaches, each associated with an OSI layer:

Repeater. A device that allows the connection of homogeneous networks at the physical layer. Repeaters are not responsible for routing or fragmentation, nor are they affected by the type of service offered (i.e. connection-oriented or connectionless).

Bridge. Operates at the data link layer and permits the interconnection of heterogenous networks. A local area network made from a number of local area networks (known as segments) is referred to as *bridged*. Two types of bridge have been considered in this chapter: transparent and source routing. Transparent bridges require less processing by the stations, but often result in the use of a nonoptimal spanning tree to ensure that packets do not circulate forever. Source routing bridges require the active involvement of the source station in the establishment of a potentially optimal route through the network.

Gateway. Like a bridge, also permits the interconnection of heterogenous networks, but operates at the network layer. Taking this approach, a network operated by different organizations supporting different protocols and addressing schemes can be interconnected. Since gateways operate at the network layer, they exist for both connection-oriented and connectionless networks. Gateways are also referred to as routers.

14.8 Exercises

1. (A) Explain both how and why a repeater connecting two bus networks permits collisions to be forwarded.

2. (A) Show, by examples, the difficulties in making synchronous mode work when connecting a pair of rings by a bridge.

3. (A) This chapter described one type of transparent bridge in which all packets are forwarded, resulting in flooding and, potentially, packets cycling forever. If a hop-count is included in each packet, does this help? How does this compare with the spanning tree and source routing algorithms? Remember, in the spanning tree algorithm the bridges are in constant communication with one another.

4. (A) An alternative method of routing is known as reverse path forwarding, in which packets are forwarded only if they are received on a link with the shortest path back to the source of the packet. Packets are forwarded on all links except the one on which they were received. Consider how reverse path forwarding can be implemented. How does each gateway "know" whether one link is a shorter path back to the source than another? Suggest a structure for the routing table. Can the reverse path ever be changed? How?

5. (A) Suggest methods whereby multicast can be supported on a bridged local area network. The solution(s) should avoid simply broadcasting each multicast message on all networks; ideally, each multicast packet will be transmitted only on networks where a member of the multicast set exists.

Can multicast be supported by the spanning tree algorithm? The source routing algorithm?

6. (A) Three approaches to initializing an intelligent bridge were suggested in this chapter. Describe the methods needed to update the station address tables if

   a. A station was removed entirely from the network.

   b. A station was removed from one network and placed on the other (keeping its physical address).

7. (A) Show the resulting spanning tree if bridge B1 in Figure 14.23 fails.

8. (B) Set up two Commkit local area networks and connect them by the bridge. To prove that the bridge actually works, transmit packets from

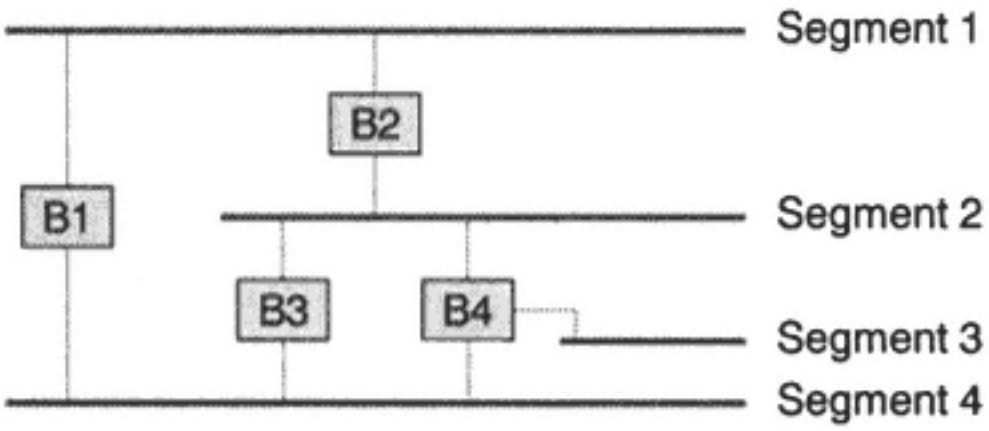

Figure 14.23
Diagram for Exercise 7.

a station on one network to a station on the other. The state of each half-bridge can be determined by querying the bridge. Now, consider the following questions:

> How long does it take for the bridge to "learn" about all the stations on the different networks?
>
> What happens to packets that are addressed to stations that do not exist on either network?
>
> What happens if two stations on different networks have the same address?

The minimum number of PCs required for this exercise is four. By alternating the nonbridge PCs on a network between the line analyzer and a network station, the traffic on the network can be monitored and traffic can be generated. It is not necessary to take the entire network down; just restart one of the nonbridge PCs with the line analyzer.

9. (B) Run a pair of networks at different speeds, one at 300 bps and the other at 9600 bps. Does everything flow smoothly? If one of the bridges crashes, explain why.

10. (C) The protocol used by `half_bridge()` does not handle lost or damaged packets. Modify `half_bridge` so that a stop-and-wait protocol handles errors in transmission.

11. (C) The routing algorithm used in `router()` works fine for a bridge connecting two networks; however, it is not general purpose and cannot be used to support more than two links. Redesign the routing algorithm so that it is general purpose and can function with anywhere from 2 to 16 links. Implement the algorithm and use the bridge keyboard to generate packets from fictitious remote networks.

12. (C) Modify the bridge software so that the Commkit wide area network can be connected to another Commkit wide area network.

13. (C) Modify the bridge software so that the Commkit wide area network can be connected to a Commkit local area network.

14. (C) What changes are necessary to the bridge software if the two interconnected local area networks support different packet sizes?

15. (B) Try running Commkit sockettes on the two networks. If the source and destination hosts are on separate networks, is the sockette packet for-

warded correctly? Are changes necessary to the networks or the bridges? Explain.

16. (C) Modify the bridge software so that Commkit sockettes will be supported using a form of IP gateway. That is, when a packet is supplied for transmission, if the destination does not exist on the local network, the packet is forwarded to the gateway. How will each host know the address of the gateway? What other changes will be required?

APPENDIXES

Appendix A
An Introduction to C

This appendix offers a brief overview of C for readers unfamiliar with the language. It covers only the language constructs used by Commkit; additional details can be found in any number of books on C or in the Turbo C manuals.

A.1
Comments

A comment begins with / * and ends with * /. Everything within it is ignored by the compiler, including any code or data structures. Comments cannot be nested but can span multiple lines.

A.2
Base Types

C supports three base types from which all other structures can be derived: integers, characters, and floating point. Only integers and characters are considered in this appendix, since Commkit does not use floating point. Any character can be used as either a character or an integer, depending upon the context.

A.3
Unnamed Constants

Commkit uses five unnamed constants:

Decimal constants. Any integers that are acceptable to the machine, such as `17, 0`, or `1027`. As shown in Table A.1, C allows decimals to be either `signed` or `unsigned`. The value of *n* in the table depends upon the context and is one of 8 (byte), 16 (word), or 32 (double word or long).

Table A.1 The range of decimal constants.

Signed		Unsigned	
Minimum	Maximum	Minimum	Maximum
2*n*1	+2*n*1 1	0	+2*n* 1

Character constants. Any alphanumeric characters enclosed in *single quotes*. for example 'A', '1', and '%'. Control characters can be declared as character constants using '\ch', where 'ch' is a lower case alphabetic character (for example, the end-of-line is delimited by '\n'). A backslash is represented as two backslashes in a row ('\ \'), while a single quote is a quote mark following a backslash ('\' ').

Octal constants. Written with the prefix `0` (zero); all digits following *must* be in the range `0` through `7`. For example, `0377` is either `-128` (if treated as a `signed` constant) or `255` (if treated as an `unsigned` constant).

Hexadecimal constants. Written with the prefix `0xhh` (where `h` denotes a hexadecimal digit, four bits long: `0` through `9, A, B, C, D, E`, and `F`). For example,`-1` can be written as `0xFFFF` (16-bit) or `0xFFFFFFFF` (32-bit).

String constants. Collections of one or more characters enclosed in *double quotes*: "`Dusty`" is an example of a string. Internally, C appends a null character ('`\0`') to the end of each string, thereby allowing an easy way to determine end-of-string. It also increases the size of the string by one byte (this is important when defining arrays of characters, as seen below).

A.4
Identifiers and Variables

An identifier is any collection of alphanumeric characters that starts with an alphabetic character. Spaces, tabs, and carriage returns are not allowed within an identifier; however, underscores '_' are allowed.

A.4.1
Variables

A variable is simply an identifier declared to be of a specific type. "Traditionally," all C variables are written in lowercase. A declaration is written as the type followed by one or more variable names (separated by commas). It is terminated by a semicolon.

Integer variables are declared as either `short` (16-bit), `int` (16-bit), or `long` (32-bit). By default, all variables are `signed`; however, the prefix `unsigned` allows the declaration of unsigned variables:

```
int alpha;
unsigned long beta;
short gamma, delta;
```

Character variables will hold one byte (8 bits) and can be used either for characters or as 8-bit integers. Characters are declared as type `char`; by default, characters are signed, although they can be explicitly declared `unsigned`:

```
char ch, data;
unsigned char subscript;
```

A.4.2
Initialization

Variables can be initialized when they are declared, for example:

```
char ch = 'X';
 int data = 3;
```

Variables that are not initialized have undefined values until an assignment takes place (see below).

A.4.3
Reserved Words

The following reserved words cannot be used as identifiers (and hence variables):

`auto`	`break`	`case`	`char`	`continue`
`default`	`do`	`double`	`else`	`extern`
`float`	`for`	`goto`	`if`	`int`
`long`	`register`	`return`	`short`	`sizeof`
`static`	`struct`	`switch`	`typedef`	`union`
`unsigned`	`void`	`while`		

The Turbo C *User's Guide* lists additional reserved words used by Turbo C.

A.5
Expressions

C supports a number of expressions and operators:

Arithmetic. The arithmetic operators found in most languages are supported by C, notably: + (addition), (subtraction), * (multiplication), and / (integer division; the decimal and fraction are discarded). The percent sign % is for integer modulus (the remainder after division). The normal precedence rules apply (i.e., multiplication, division, and modulus have higher priority than addition and subtraction). Parenthesis can be used to change the order of evaluation, for example:

```
2 + 3       * 4       /* = 14 */
(2 + 3)       * 4     /* = 20 */
```

Note that arithmetic operations *can* be applied to characters. For example, 'A' + 1 gives 'B' (since the character after 'A' in the ASCII collating sequence is 'B').

Relational. Relational operators allow the comparison of two expressions. C's six relational operators are shown in Table A.2. All relational expressions evaluate to either 0 (false) or 1 (true). There are no built-in constants TRUE or FALSE (as in Pascal).

Booleans are built out of integers (or characters). Zero is false, while anything else (typically 1) is taken as true.

Shift. There are two shift operators: << for left shift and >> for right shift. The lefthand expression is the variable (or expression) to be shifted, while the righthand expression is the amount by which the lefthand expression is to be shifted. For example, x << 3 shifts the variable x to the left by 3 (this is equivalent to multiplying by 8).

Bitwise. Bits can be set and cleared using the bit operators shown in Table A.3. For example, to mask the lower eight bits of an integer i, the bitwise "and" operator can be used: i & 0xff.

Logical. There are two logical operators: && (and) and || (or). They have the lowest precedence and are evaluated from left to right. A zero-valued expression is interpreted as false, while all other values are taken as true. Logical expressions are evaluated from left to right until there is no longer any need to continue evaluating. For example, a logical expression consisting of *or* operators (||) is evaluated until the first true (i.e., nonzero) expression is found, at which point the entire logical expression is taken to be true. Similarly, an expression using the *and* operator (&&) is evaluated until the end of expression is reached *or* one of the expressions is found to be false.

Conditional operator. The conditional operator ? allows an if-then-elselike construct to be embedded within an expression:

Table A.2 Relational operators.

Symbol	Operation
ex1 = = *ex2*	Equals
exl ! = *ex2*	Not equals
ex1 < *ex2*	Less than
exl > *ex2*	Greater than
ex1 <= *ex2*	Less than or equal
ex1 >= *ex2*	Greater than or equal

Table A.3 Bitwise operators.

Symbol	Operation
|	Bitwise "or"
&	Bitwise "and"
~	Exclusive-or

(*expression*) *? true-part : false-part*

For example, a problem may require `6` to be added to x if the `total` equals `3`; otherwise, `4` is to be added to x; this can be written as follows:

```
x + (total == 3) ? 6 : 4
```

A.6
The Assignment Statement

The assignment statement is defined as a left value (*1value*) being assigned the result of a right value (*rvalue*). It is written as *1value = rvalue*. The *1value* is always a memory location and the *rvalue* an expression. Unless otherwise indicated, the statement is terminated with a semicolon (;). C supports little or no checking when dealing with variables of the base type, for example:

```
int a;
char b;

a = 'X';        /* Assigning a character to an
                   integer */
b = a + 1;      /* Storing an integer into a
                   character variable */
```

Multiple assignments are allowed:

```
a = b = c = 10;
```

Beware of seemingly innocent typos such as

```
a = b == c = 10;
```

in which `c` is assigned the value `10`, then `b` is compared with `c`, and the result of the comparison (`0` or `1`) is assigned to `a`.

A.6.1
Variations

C offers a number of shorthand notations for the assignment statement:

Statements such as `x = x` <*op*> `y` can be abbreviated to `x` <*op*>`= y`. The <*op*> can be one of +, , *, /, %, <<, >>, &, |, or ^ . For example, the statement `a = a + b` can be abbreviated to `a += b`.

Increments (such as `a = a + 1`), can be abbreviated to `a++` or `++a`. There *is* a difference between `a++` and `++a`; `a++` indicates that the value of `a` is to be used in the expression; once taken, `a` is to be incremented. On the other hand, `++a` indicates that a is to be incremented and then its value is to be used in the expression. Similarly, decrements can be abbreviated to `a` or `a`. The same rules apply to the positioning of the signs.

It is important to note that an assignment can take place *anywhere*. For example, the following statement:

```
a = (b = 3) * 2;
```

results in `b` being assigned the value `3`, while `a` is assigned the value `6` (i.e., the value of `b` after the assignment times `2`).

A.7
Selection

C supports two selection statements, one conditional and the other a multiway branch.

A.7.1
Compound Statements

Compound statements are groups of zero or more statements enclosed in braces `({. . .})`; note that *all* statements must end with a semicolon:

{ *Statementl*;
`/* More statements */`
StatementN;
}

The compound statement does *not* end with a semicolon.

A.7.2
The if Statement

The `if` statement is written as

`if` (*Expression*)
Statement1;
`else`
Statement2;

The *Expression* (see above) is evaluated; a non-zero result causes *Statement1* to be executed; otherwise, *Statement2* is executed. If the `else`*Statement2* construct is omitted, the result is an if-then statement. Note that *Statement1* and *Statement2* can both be compound statements (remember that compound statements cannot be followed by a semicolon).

The following code fragment illustrates an *if* statement: should a equal 'X' or `c` be less than `2`, data is assigned the value `5`; otherwise, `data` is cleared and `a` is assigned 'Z'.

```
if (a ==  'X' || c < 2)
      data = 5;
else
{
      /* a != 'X' and c >= 2 */
      data = 0;
      a = 'Z';
}
```

A.7.3
The switch Statement

The multiway branch is known as the *switch* statement; it is normally written in the following form:

`switch`(*Expression*)
`{`
`case` *Constant*:
Statement(s);
`break`;

`case` *Constant*:
Statement(s);
`break`;

```
/* Other statements */
```

`default:`
Statement(s);
`}`

The *Expression* is evaluated to an integer value; control is passed to the `case` label (a constant), which matches the value of the *Expression*. The *Statement(s)* following the label are then executed. If a section of code is

to be associated with a number of different values of the *Expression*. each *Constant* must be associated with its own `case` label; for example:

```
switch (ch)
 {
 case 'A':
 case 'a':
   /* Statements */
   break;

 case 'B':
 case 'b':
   /* Statements */
   break;

 /* Other 'case' labels and statements */
 }
```

Once the set of statements associated with the *Expression* has been evaluated, control can be passed outside of the `switch` statement using the `break` statement. It is possible to branch into the middle of a series of statements simply by placing the `case` label above the first statement associated with the `case` label:

```
switch (ch)
{
case 'A':
     ch = 'a';
case 'a':
     /* Statements */
     break;

     /* Other 'case' labels and statements */
}
```

If the value of the *Expression* does not match any of the `case` labels, control passes to the statements that follow the label `default`:. If there is no `default`, control passes to the first statement following the closing brace of the `switch`.

A.8
Iteration

C supports three structured iteration statements as well as a `goto` statement.

A.8.1
The while Statement

The `while` statement is a pretest, nondeterministic loop structure, written in the form

`while` (*Expression*)
Statement;

The *Expression* is evaluated; if it is non-zero, the *Statement* is executed. The cycle is repeated as long as the result of the expression is non-zero. The *Statement* can be compound, for example:

```
count = 0;
while (count < 10)
{
   /* Other statements */
   count++;
}
```

Often the loop can proceed backwards, producing some interesting software:

```
count = 10;
while (count--)
{
      /* Statements */
}
```

The loop will be entered with the final value of `count` being `0`; the next iteration will determine that `count` has a `0` value and the loop will terminate.

An infinite loop can be written by setting the *Expression* to `1: while (1)`.

A.8.2
The do. .while Statement

The `do..while` statement is a posttest, nondeterministic loop, written in the form

`do`
Statement;
`while` (*Expression*);

The *Statement* (which can be compound) is executed before the *Expression* is evaluated. The cycle continues as long as the *Expression* produces a non-zero result. Multiple statements must be written as a compound statement.

A.8.3
The for Statement

The `for` statement allows the construction of deterministic loops (i.e., loops with a known initial condition, final condition, and increment). The format of the `for` statement is as follows:

`for` (*Expression1; Expression2; Expression3*)
Statement;

where *Expression1* is the initial condition (typically an assignment), *Expression2* is the termination condition, and *Expression3* is the increment. For example, to count from `0` to `10`, a `for` loop can be written as:

```
for (i=0; i<=10; i++)
{
   /* Statements */
 }
```

Note that the `for` loop is equivalent to

```
expression1;
while (expression2)
{
     /* Statements */
     expression3;
}
```

Finally, any or all of the *expression*s may be omitted. For example, the following set of statements are performed "forever"

```
for(;;)
 {
   /* Statements */
 }
```

A.8.4
The goto Statement

An unconditional transfer of control can be achieved using the `goto` statement, which is written with an identifier (a *label*), for example:

```
goto done;
```

The label must be within the same function (see below) as the `goto`, is ter-

minated with a colon (not a semicolon), and can branch forward or backward over any number of nested loops:

```
while(1)
{
     /* Statements */
     if (data == 'X') goto done;
     /* Statements */
}
done:
/* Statements */
```

Note that `goto` is different from `break` in that it can branch anywhere within a function. However, `break` is more structured, since control passes to the first statement beyond the end of the block in which the `break` is written. The `continue` statement passes control to the end of the block in which the `continue` is written.

For example, if the statement `goto done` is replaced by `break`, execution resumes with the first statement outside the `while` loop. However, if `continue` replaces `goto done`, the statements between the `continue` and the closing brace are ignored, with execution resuming at the start of the loop (i.e., the `while`).

A.9
Aggregate Data Types

C allows complex data types (notably arrays, structures, and unions) to be constructed out of the three base types. Additionally, pointers to the base or aggregate types can be constructed.

A.9.1
Arrays

An array is a data structure consisting of one of more elements sharing a common type and name (an identifier). It is declared by specifying its type, name, and dimension. For example, to declare an array of 10 integers, the following declaration can be used:

```
int data[10];
```

An individual element in the array is accessed using a subscript enclosed in square brackets. Subscripts are integers (or characters) and must be in the range `0` through `N-1` (where `N` is the size of the array). For example, the array `data` can be set to `0` using a `for` loop:

```
for (i=0; i<10; i++)
     data[i] = 0;
```

Arrays can also be initialized when they are declared:

```
int data[10] = {0, 0, 0, 0, 0, 0, 0, 0, 0, 0};
```

An array can have a maximum of two dimensions, where each dimension is separated by a comma.

A string is simply an array of `chars`. For example:

```
char name[10];
```

Text strings cannot be assigned directly to string variables in an assignment statement, although individual characters may be assigned to each array element. However, C has many string manipulation routines that can be used to access, compare, and manipulate strings.

Strings can be initialized at compile time in much the same way integer arrays are handled (note that the braces are omitted):

```
char name[10] = "Your name";
```

In the above example, the array name is assigned the nine characters of the string `"Your name"`. A tenth character (the null character) is added at the end of the string. To avoid counting each character in a string, C allows a shorthand notation for character string initialization:

```
char name[] = "Your name";
```

A.9.2
Structures

Separate but related data structures can be placed in a single larger data structure known as a `struct`. The basic format of a structure is

```
struct
{
field(s);
}
```

The structure consists of one or more *field(s)*, where a *field* is a data structure declaration. For example, a person's birthday consisting of a day, month, and year consists of related items that can be grouped into a `struct`:

```
struct
{
int day;
int month;
int year;
}
```

A structure can be used to declare a new data structure or a new data type (or both); the above example is incorrect in that the structure has not declared a new data structure or new data type.

A data structure is declared with its name following the closing brace. For example, the data structure `my_birthday` with the fields day, month, and year can be declared as follows (note that the structure ends with a semicolon):

```
struct
{
int day;
int month;
int year;
} my_birthday;
```

To declare a new data type, its name is entered after the word struct and before the `{`for example (note that the structure must be terminated with a semicolon after the closing bracket):

```
struct birthday
{
int day;
int month;
int year;
};
```

Structures can be declared within other structures.

The rules for declaring a data structure of type `struct` are the same as for any other declaration: the name of the type (for example, `struct birthday`) must be followed by one or more *identifiers*. separated by commas and terminated with a semicolon:

```
struct birthday evans_birthday;
struct birthday the_cats_birthday;
```

The individual fields within the structure are accessed by specifying the structure's name (i.e., the *identifier*), followed by a `.', followed by the name of the field (note that this can be recursive if structures within structures are declared). Structures can also be initialized at compile time. For example:

```
struct birthday evans_birthday = {18, 5, 1978};
struct birthday the_cats_birthday;

the_cats_birthday . day   = 1;
the_cats_birthday . month = 4;
the_cats_birthday . year = 1990;
```

The individual fields within the structure can be manipulated on the basis of their type.

Structures can be declared as arrays and accessed using subscripts:

```
struct birthday cat_family[5];
int i;

for (i=0; i<5; i++)
{
     cat_family[i] . day   = 0;
     cat_family[i] . month = 0;
     cat_family[i] . year  = 0;
}
```

A.9.3
Unions

Data structures can share the same memory locations using a `union`. A `union` is declared and accessed in the same way as a structure, the difference being that each *field* entry in a `union` refers to the same memory location. For example, the following `union` declaration allows a 32-bit location to be accessed as four bytes, two words, or one long word:

```
union memloc
{
char byte[4];
int word[2];
long double_word;
};

union memloc x;
```

The variable `x` refers to a single 32-bit location and can be visualized as follows:

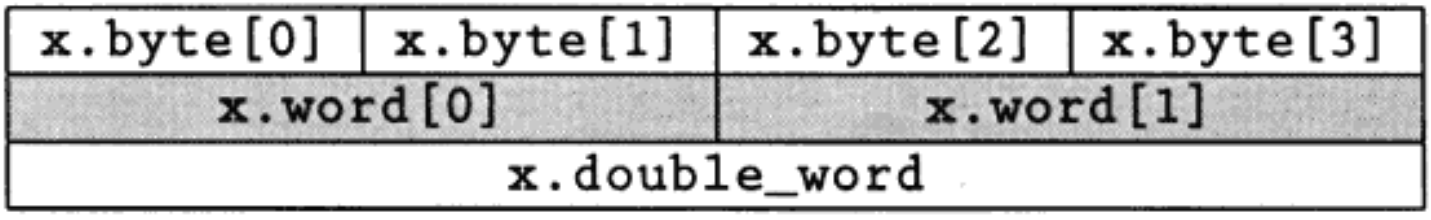

A.9.4
Pointers

All data structures are associated with an address. C allows the program to access a data structure through its name or its address. The address is obtained by placing an & before the data structure name. For example, the address of an integer `x` can be obtained by writing & before `x`.

A pointer is declared as one to a specific type. A pointer to an integer is thus declared as

```
int *ptr;
```

Pointers are assigned values (usually addresses, although not a necessity because of C's lax type checking) using an assignment statement. To refer to the location indicated by the pointer requires placing an * in front of the pointer's name.

A typical, contrived, example of how a pointer functions is

```
int *ptr;      /* A pointer to an integer */
int data, ans; /* Two integers */

ptr = &data;   /* 'ptr' now contains the address
                  of 'data' */
*ptr = 7;      /* 'data' now has a value of 7 */
ans = *ptr;    /* 'ans' takes the value of the
                  location */
               /*  pointed to by 'ptr' (i.e. 7) */
```

Pointers can point to array elements as long as the types agree, for example:

```
char *cptr;          /* A pointer to a character */
char array[10];      /* A string of 10 characters */

cptr = &array[2];    /* 'cptr' points to the 3rd
                        element in 'array' */
 *cptr = 'S';        /* 'array[2]' now contains 'S' */
```

Pointers can also be incremented and decremented. For example, to initialize `array` to `?', one can write:

```
cptr = &array[0]; /* or simply 'cptr = &array' */
i = 0;
while (i < 10)
{
     *cptr++ = '?'; /* Assign '?' then increment
                       'cptr' */
     i++;
}
```

Pointers can also point to structures and unions. A structure (or union) pointer is declared to be to the specific structure. In reference to a *field* within the structure, the pointer name is followed by `->` and finally the field name:

```
struct birthday evans;
struct birthday *guess; /* Pointer to struct
                        'birthday' */

evans . day = 18;
evans . month = 5;
evans . year = 1978;

guess = &evans;           /* Address of struct
                          'evans' */
guess-> day-= 5;        /* Decrement 'day' by 5 */
guess-> month = 2;       /* Change 'month' to 2 */

if (guess-> year > 1950)
     guess-> year = 1949;   /* Change year to 1949 */

/* 'evans' now contains 13 (day), 2 (month),
    1949 (year) */
```

A.10
Functions

A C program consists of one or more functions. All functions have the same format, notably:

Result-Type Function-Name (*Parameter-Declarations*)
{
Function-Body
}

The *Result-Type* can be a base (`int, char, long, unsigned`, etc.) or aggregate. However, if an aggregate type is being returned, it should be returned as an address since the function returns at most a 16- or 32-bit integer value. A *Result-Type* of `void` indicates that nothing is to be returned, meaning that the function is essentially a procedure. The *Function-Name* is a valid identifier name, while the *Function-Body* is enclosed in braces: `{. . .}`. The *Function-Body* consists of (local) variable declarations as well as executable statements.

If the *Result-Type* is omitted, the function is assumed to return an integer. The *Parameter-Declarations* are optional: if they are omitted, the parenthe-

sis must follow the *Function-Name*. For example, the function `ex1()` is an integer function:

```
ex1()
{
/* Statements */
}
```

All parameters are considered local to the function and, when listed, must be separated by commas. For example, the following function is of type `int`, with three parameters (`arg1` is an integer, `arg2` is a character, and `arg3` is a pointer to an integer):

```
int example(int arg1, char arg2, int *arg3)
{
/* Statements */
}
```

A value can be returned from a function using the `return` statement. For example, the following function returns the larger of two integers:

```
int largest(int data1, data2)
{
return (data1 > data2) ? data1 : data2;
}
```

A function is called by writing the *Function-Name* followed by the arguments associated with the function. For example, to find the largest of two numbers, `num1` and `num2`, one can write

```
answer = largest(num1, num2);
```

It is possible to ignore the return value by *casting* the function to void:

```
(void) largest(num1, num2);
```

Unless otherwise specified, all parameters are call by value, meaning that, whatever changes take place to the parameter in the function, the corresponding argument remains unchanged. Should it be necessary to have the function change the value of the argument, C allows the arguments to be passed by reference.

A call-by-reference parameter requires the address of the data structure to be the argument; the corresponding parameter in the *Parameter-Declarations* must be a pointer to the specified type. Structures *must* be passed by reference. Consider the following example:

```
void ex2(struct birthday *bptr, int *iptr)
{
bptr -> day = 26;
bptr -> month = 8;
bptr -> year = 1954;
*iptr = 123;
}

void call_ex()
{
struct birthday jaws;
int dusty;

ex2(&jaws, &dusty);

/* jaws: 26 (day), 8 (month), 1954 (year),
   and dusty: 123 */
}
```

All variables declared within a function are local to it. Global variables are declared outside of functions; they are global to all functions. Aggregate types can be declared globally as well. Since C programs can be developed in a number of different files, global data structures (common to a number of separately compiled functions) can be declared as externals using the `extern` type. For example, assuming that a number of separately compiled functions share a common data structure, `cookie`, of type `struct birthday`, one file requires the declaration `struct birthday cookie` (to reserve the memory location), while the other files contain the declaration `extern struct birthday cookie`. The linker resolves any addressing problems.

An alternative to declaring a global variable that is used by a single function is to declare a local static variable. The `static` variable retains its value between calls of the function, whereas all other local variables are automatic in that they are created on the stack for the duration of the function's call. A `static` variable can be initialized at its declaration:

```
int example()
{
static char data = 'X';
/* Statements */
}
```

The entry point from the operating system into the program must be a function with the name `main()`. This function can have two parameters, the first indicating the number of items entered on the command line when the program is loaded, and the second an array of pointers to each word

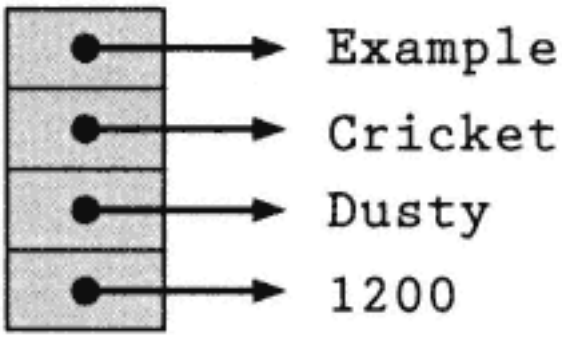

Figure A.1
`Argv` is a list of character pointers.

(assumed to be a character string) entered on the command line. These two parameters are given the names `argc` and `argv`, respectively:

```
main(int argc, char *argv[])
{
/* Statements */
}
```

For example, if an executable program `example` has three arguments entered on the command line as follows:

```
C:\> example cricket dusty 1200
```

then the value of `argc` is `4` (there are four "words" entered on the command line), and `argv` is an array of string pointers; the structure is shown in Figure A.1.

Any of the strings can be accessed; for example, to access `dusty`, one refers to the third element of `argv, argv [2]`.

Some general points about functions:

Functions cannot be local to other functions; that is, all functions are global.

Forward references to functions are allowed; however, the functions are assumed to return integers. Function headers (the *Result-Type*. the *FunctionName*. and the *Parameter-Declaration*) can be defined as global statements at the start of the file or even externally.

By default, the compiler passes all strings by reference.

Unless the *Result-Type* is specified, the function is assumed to be of type `int` (returning an integer value).

A.11
Compiler Directives

C supports a number of compiler directives that instruct the compiler to perform an action that need not result in the generation of code. Two compiler directives used by Commkit are `#define` and `#ifdef`.

The `#define` compiler directive instructs the compiler to store a symbol and a value in the symbol table. A common use of `#define` is to declare *named constants*. for example:

```
#define TRUE      1
#define FALSE     0
#define LIMIT     25
#define MASK      Oxff
#define VALUE     'w'

main()
{
char data[LIMIT]; /* 'data' is an array of size
                  LIMIT */

if (data[3] & MASK == VALUE)
  then data[3] = 0;
}
```

It is common to write all defined symbols in uppercase to distinguish them from other data structures. Defined symbols are *not* variables; they cannot be an 1*value*. nor are they associated with an address.

The #define directive can be used for more than simply defining named constantsit can define entire expressions or statements. For example:

```
#define FOREVER      for(;;)
#define DOUBLE_X     x *= 2;
```

Whenever a defined symbol is written, the compiler expands it into whatever it is defined as. For example, whenever `FOREVER` is encountered, the compiler actually compiles `for (;;)`.

Arguments can be passed to compiler directives as well. For example, to allow any value to be doubled (instead of `x` as in the previous example):

```
#define DOUBLE(value) value *= 2;

main()
{
int x, count;

DOUBLE(x)     /* Compiler produces x *= 2;      */
DOUBLE(count)  /* Compiler produces count *= 2; */
}
```

Multiple arguments are allowed, although the exact number depends upon the compiler.

Conditional compilation is possible using the `#ifdef` compiler directive in conjunction with the `#define` directive. Conditional compilation permits the programmer to instruct the compiler to generate code under certain conditions (for example, when searching for an error).

For example, to track down an error it is possible to plant diagnostic statements throughout a program. Once the error is found, all the diagnostic statements may be removed (although existing code may be damaged). An alternative is to leave the diagnostic software in the program but to associate the diagnostic statements with directives that tell the compiler when to include the diagnostics:

```
void a_procedure()
{
   /* Statements */
#ifdef DEBUG
   /* Diagnostic statements */
#endif
   /* Statements */
}
```

The compiler will include the diagnostic statements between the `#ifdef` and `#endif` if `DEBUG` has been defined; otherwise, the statements are left out of the compilation. `DEBUG` can be defined simply by writing `#define DEBUG` (there is no need to associate a value with `DEBUG`, the compiler simply marks it as defined).

A.12 Software Management

There are two types of source code file: *source* (with a `.c` extension), that is, C programs that can be compiled; and *header* (indicated by the `.h` extension), containing definitions and data structures. All of the source files have an equivalent *object* version (each with an extension `.obj`) consisting of executable software that must be linked (with other object files) to create an executable file (with an `.exe` extension).

To minimize the amount of compiling required each time a change is made to a source file, a software management tool known as `make` is supplied with Turbo C. The `make` utility controls the recompilation of files by reading the commands specified in the `makefile`. The `makefile` contains a list of *dependencies* that specify the source files that must be recompiled and relinked after a change is made. For example, consider the steps in the creation of `ipc.exe`:

```
ipc.exe: ipc.obj commkit.obj srcalls.obj
    tlink \tc\lib\c0s ipc commkit srcalls, ipc, , \
    \tc\lib\emu \tc\lib\maths \tc\lib\cs
```

```
ipc.obj: ipc.c general.h ascii.h devices.h
    tcc -ms -c ipc
```

The file `ipc.obj` is dependent upon `ipc.c`: should a change occur to `ipc.c` (i.e., if the time and date of `ipc.obj` are earlier than those of `ipc.c`), `ipc.c` is recompiled using `tcc`, the Turbo C (and C++) compiler, producing a new copy of `ipc.obj`. The executable version of `ipc.c (ipc.exe)` is also specified as a dependency; this time the Turbo linker `tlink` is called to link the object files `ipc.obj`, `commkit.obj`, and `srcalls.obj` to create `ipc.exe`.

The instructions in the `makefile` are processed by the `make` utility by typing

```
C:\> make
```

When the `make` utility finds a file that must be recompiled or relinked, the specific line in the `makefile` is displayed. If all files are found to be up to date, `make` returns to the MS-DOS prompt.

A specific file can be processed by the `make` utility by typing the file name after `make`. For example, to check whether `ipc.exe` is up to date, one types

```
C:\> make ipc.exe
```

In some cases it is necessary to remake an executable module without having modified the original source module. For example, all object and executable modules supplied with Commkit were created using Turbo C++; recompiling all the source modules with Turbo C means having to change the date on each module using a tool such as an editor. Fortunately, the `touch` utility (supplied with Turbo C and C++) can change a file's time and date to the present; if `touch` is followed by a `make`, the associated files will be recompiled. As an example, to force the recompilation of `ipc.c`, one can type

```
C:\> touch ipc.c
C:\> make ipc.exe
```

Wildcards work with `touch;` a filename of `*.c` will change the time and date of all `*.c` modules to the present.

Appendix B
Running and Testing Commkit

One of the biggest sources of frustration for any programmer is to be stuck with software that doesn't function as expected. This is as true for communications software as it is for any other type. This appendix suggests a number of techniques that can be used to help find and correct errors.

B.1
General Hints

There are no hard and fast rules for modifying software and checking for errors. However, the following points may help reduce some of the frustration of working with communication software developed by someone else:

Use the line analyzer to monitor the communication channel. The line analyzer removes as much of the guesswork as possible about what is occurring on the channel. Knowing what is happening "down there," often makes it easier to determine what is happening "up top."

Timing is often critical when testing communication software: avoid writing software that "assumes" a specific line speed, because it can break down at different speeds. A good example of this can be found in ring network software: the ring monitor is designed to wait five seconds without sensing a token; after this time all traffic is removed from the network and a new token is placed on the ring. The software works correctly at all speeds except 50 bps. Since that speed is roughly 5 characters per second, a message of more than about 25 characters will cause the monitor to reset the network. (There is an exercise in Chapter 12 for this specific problem!)

Don't make too many changes at once: make a few and then exercise the software. If an error occurs, one can be confident that it has been caused by the new software.

It is also wise to avoid making changes to the line analyzer at the same time that other software is being changed. For example, an error occurring after modifying both the line analyzer and the point-to-point software can lead to untold hours of frustration, because there are two possible sources of the error.

Check the wiring that interconnects the various PCs before blaming the software. A good example is the ring network, which requires a specific ordering of connections (i.e., serial port 1 to serial port 2). If the wiring is wrong, the software can't be expected to function properly.

The message buffer queue can "suddenly" empty if messages are being sent but never received, resulting in the queue of pending messages being displayed on the PC and the software terminating. The first byte displayed after dump is the process (i.e., queue) number: this will indicate which process is not reading its queue.

Remember, the background processes are only invoked when the foreground process performs a `recv()`. The foreground process can cause the message buffer queue to empty if it sends vast numbers of messages to a background process while never pausing to perform a `recv()`.

When testing network software, put sequence numbers into the messages being sent; the sequence numbers can determine whether messages have been lost, corrupted, or are simply hiding somewhere waiting to be transmitted. The line analyzer displays all binary information as a byte using the PC's character set; as the sequence number increases, the sequence number displayed by the line analyzer will progress through the character set table.

If possible, avoid using `printf()` anywhere other than in the foreground process. If the implementation of `printf()` is not re-entrant, results can be unpredictable should an interrupt handler access `printf()` at the same time the foreground process does. Use `message()` or `display()` in interrupt handlers.

Always take backup copies of working software; there's no enjoyment in retyping a set of software from a week-old listing.

Although most compiler warning messages can safely be ignored, they often provide clues as to why things aren't working.

If all else fails, Commkit recognizes CTRL-ALT-DEL as a catastrophic abort and will return control to MS-DOS.

B.2
A Final Thought

Remember, testing indicates only the presence of errors, not their absence. In other words, that a test has run for an hour and an error hasn't occurred doesn't mean that the software is error-free. Design the tests to exercise as many parts of the software as practical (i.e., both the true and false parts of an `if` statement and the different `case` options in a `switch` statement).

Appendix C
State Diagrams

Many communication systems and protocols (including some in this text) are described in terms of state diagrams or state machines. State diagrams permit an unambiguous, visual description of a system's individual events and their related actions. This appendix presents an overview of state diagrams and includes an example of the design and implementation of a state diagram and its associated application.

C.1 Concepts

Any task that can be broken into a series of distinct steps (or states) can be defined in terms of a state diagram. Control remains in a state until an event occurs, which causes a transition to a new state. For example, the task of crossing a road at a crosswalk can be broken into two states:

1. Wait for a green light.

2. Cross the road.

The event that causes the state transition (from waiting for a green light to crossing the road) is the availability of a green light. In other words, when the light turns green, the road can be crossed.

A state diagram is a visual representation of a task's various states and events. Formally, it is a weighted, directed graph, with each vertex corresponding to a state and each directed edge representing a transition. At a minimum, a state diagram consists of three parts: the state (typically a labeled box), the transition between states (typically an arrow joining the two boxes), and a description of the event that caused the transition. Figure C.1 is the state diagram for the street-crossing task.

Each event (or condition) that causes a transition to a new state is associated with an action or output. In a state diagram, this is drawn as follows:

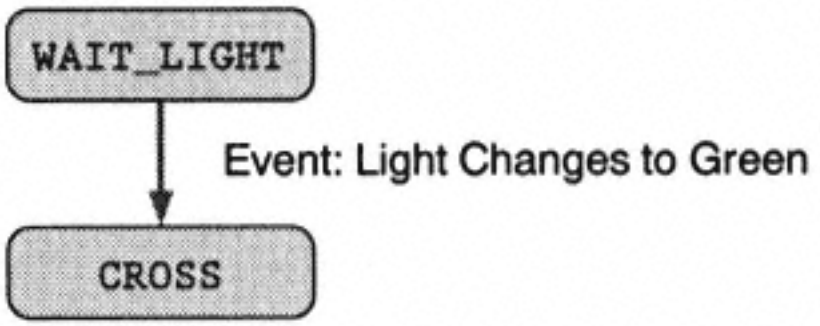

Figure C.1
State diagram for crossing a street.

Condition
Action

Note that in some cases, such as the state diagram for crossing the road, there is no action associated with the condition.

As an example of a state diagram with both conditions and actions associated with a transition, consider the traffic light's states:

1. Show the red light until the timer expires, at which point, turn the light green and start the timer for 50 seconds.

2. Show the green light until the timer expires, at which point, turn the light amber and start the timer for 10 seconds.

3. Show the amber light until the timer expires, at which point, turn the light red and start the timer for 60 seconds.

In this task, control remains in each state until the timer expires, at which point the light changes color and the timer is assigned a new value. A possible state diagram for the traffic light is shown in Figure C.2. Note the use of the initialization step, which causes the timer to run for 60 seconds and the light to turn red.

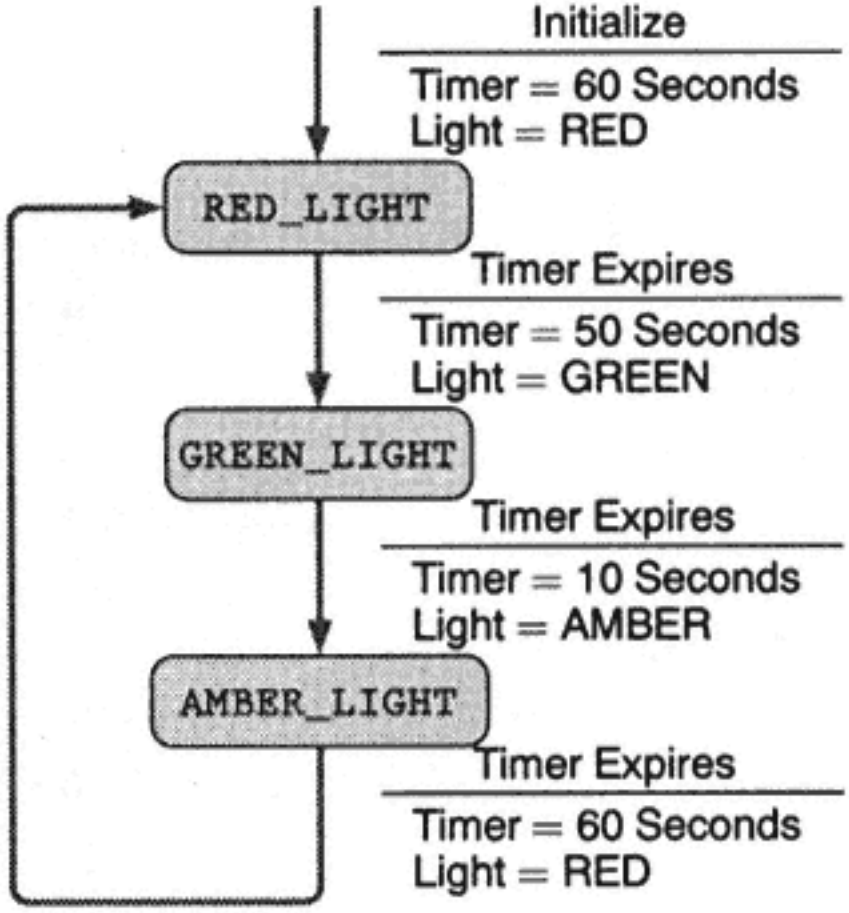

Figure C.2
State diagram for a traffic light.

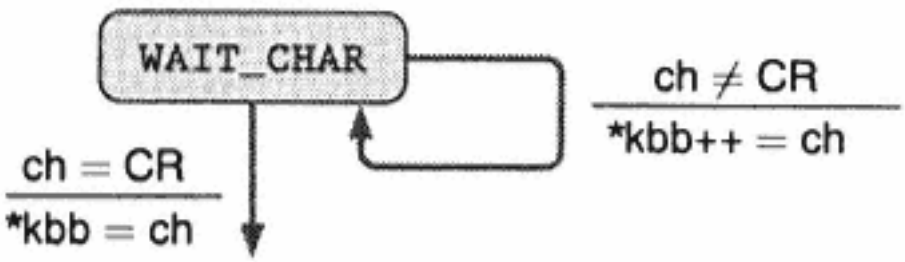

Figure C.3
State diagram for storing keystrokes.

A state can be associated with transitions to more than one state. In these situations, different conditions are applied to determine the new state. Furthermore, it is possible for a condition to lead back to the same state; for example, control may stay within a state until a certain number of events have occurred. The example in Figure C.3 shows how incoming keystrokes can be monitored; if a carriage return is detected, control passes to the next state; otherwise, the character is stored in the keyboard buffer. Note that the event in this example is the arrival of a character that is subsequently examined.

C.2 Implementation

Once a state diagram has been developed for a task, it can be implemented. Implementation is most easily achieved using a `switch` statement; the `switch` expression is the new state, and the different states are listed as `case` labels.

In the state diagram, control remains in the state until a condition arises that causes a transition to a new state. In a system of multiple processes, it may not be possible to remain physically in the state while waiting for the event to occur; instead, the process must relinquish the CPU until that time. For example, in Commkit, when a task has completed the actions associated with a state it calls `recv()`.

As an example of implementing a state diagram, consider the traffic light task. In this case, the state diagram consists of four transitions and three states. The first transition, initialization, consists of starting the timer and turning the light on and then entering the `RED_LIGHT` state. The task, running as long as `power_on` remains `TRUE`, waits until the timer expires, at which point the traffic light state changes to its new state. The traffic light state diagram can be implemented as follows:

```
/* Initialization */
start_timer(60);
light_on(RED);
state = RED_LIGHT;
while(power_on)
```

```
{
   wait_for_timer();
   /* At this point, the timer has expired */
   switch (state)
   {
   case RED_LIGHT:      /* Light is red,
                           turn to green */
      start_timer(50);
      light_on(GREEN);
      state = GREEN_LIGHT;
      break;
   case GREEN_LIGHT:    /* Light is green,
                           turn to amber */
      start_timer(10);
      light_on(AMBER);
      state = AMBER_LIGHT;
      break;
   case AMBER_LIGHT:    /* Light is amber,
                           turn to red */
      start_timer(60);
      light_on(RED);
      state = RED_LIGHT;
      break;
   }
}
```

When several possible conditions can lead from a state, it is necessary to select the new state by testing the various conditions. For example, in the keystroke state diagram the state changes only when a carriage return is detected:

```
case WAIT_CHAR:
   if (ch != CR)
       *kbb++ = ch;
   else
   {
       *kbb = NUL;
       state = END_OF_LINE;
}
```

Appendix D
Character Codes

In its most basic form, all information within a computer is represented as a data structure consisting of one or more bits. The "fundamental" data structure that most computer manufacturers refer to when describing their equipment is the byte (for example, many PCs are sold with 640 kilo*bytes* of storage).

The byte itself can be interpreted in a number of ways, for example, as an integer or as a character. Various character codes have been developed over the past century that specify a standard to which manufacturers conform when mapping the bits making up the byte into a "character." The character need not be printable; some characters are treated as control codes used in the transmission of information.

This appendix considers three such character codes: Baudot, ASCII, and EBCDIC.

D.1
Five-Bit Baudot

Baudot is a 5-bit code, in which five bits represent a single character (in many documents, the word *unit* replaces the word *bit*). It is named after a Frenchman, Emile Baudot, an early pioneer in telegraphy who developed the code in the 1870s. Teletypewriter services such as Western Union use equipment that recognizes the Baudot code for the transmission and reception of Telex messages. The 8250 UART, used by the PC, can be configured to accept 5-bit data.

At first it may appear that a 5-bit code is of little use, since a maximum of only 32 (i.e., 25) character code values are possible. However, by defining "shift" characters, the number of possible character values *doubles* to 64. There are two shift characters:

LTRS. The incoming data is to be interpreted as letters.
FIGS. The incoming data is to be interpreted as
figures.

Table D.1 The Western Union 5-bit Baudot code.

	Letters	Figures
00000	Blank	Blank
00001	E	3
00010	*LF*	*LF*
00011	A	-
00100	Space	Space
00101	S	'
00110	I	8
00111	U	7
01000	*CR*	*CR*
01001	D	*WRU*
01010	R	4
01011	J	*BELL*
01100	N	,
01101	F	$
01110	C	:
01111	K	(

	Letters	Figures
10000	T	5
10001	Z	"
10010	L	)
10011	W	2
10100	H	#
10101	Y	6
10110	P	0
10111	Q	1
11000	O	9
11001	B	?
11010	G	&
11011	*FIGS*	*FIGS*
11100	M	.
11101	X	/
11110	V	;
11111	*LTRS*	*LTRS*

Although there are many implementations of the Baudot code, all characters received as *LTRS* are interpreted in the same way; however, *FIGS* may be unique to a specific application. The ITU International Alphabet Number 2 (or IA2) is a 5-bit Baudot code used for telex communications.

Table D.1 shows the Western Union telex 5-bit code (the "Letters" column is common to *all* 5-bit Baudot codes, not only Western Union).

There are four special characters:

WRU is a shorthand notation for "who are you."
BELL causes a bell to ring on the receiving teletype.
CR is a carriage return.
LF is a line feed.

As an example, if the bit pattern `01110` is received, it is taken either as a *C* in *LTRS* mode or as a ":" in *FIGS* mode. The *LTRS* and *FIGS* characters are transmitted only when the user shifts from letters to figures or vice versa.

D.2 Seven-Bit ASCII

Five-bit codes such as Baudot are restrictive in that only uppercase characters are handled. The interpretation of each character depends upon whether a shift character is properly received, and there is no room for control characters.

In light of this limitation, a number of organizations together developed the 7-bit ASCII (American Standard Code for Information Interchange) code, formally known as ANSI (American National Standards Institute) standard X3.4-1977. ASCII supports 128 possible characters (i.e., 27): upper- and lowercase characters, numbers, special symbols, and control characters. The ITU equivalent of ASCII is known as ITU International Alphabet Number 5 (or IA5).

The 7-bit ASCII code is shown in Table D.2. To read the bit pattern associated with a specific character, find the character in the table, and concatenate the three bits at the top of the character's column with the four bits to the left of the character. For example, the letter S is the bit pattern `101` concatenated with `0011`, or `1010011`. The ASCII control characters are defined in Table D.3.

Originally, 7-bit ASCII was developed for machines supporting 7-bit bytes (actually eight bits in total: the eighth bit of a character was treated as a parity bit). However, since most bytes occupy eight bits within the machine, many manufacturers have extended the 7-bit ASCII character set to support a further 128 characters. For example, the PC treats bytes with values greater than 127 as special graphic characters, generated using the ALTMODE key.

Table D.2 The 7-Bit ASCII Character Code

	000	001	010	011	100	101	110	111
0000	NUL	DLE	SP	0	@	P	`	p
0001	SOH	DC1	!	1	A	Q	a	q
0010	STX	DC2	"	2	B	R	b	r
0011	ETX	DC3	#	3	C	S	c	s
0100	EOT	DC4	$	4	D	T	d	t
0101	ENQ	NAK	%	5	E	U	e	u
0110	ACK	SYN	&	6	F	V	f	v
0111	BEL	ETB	'	7	G	W	g	w
1000	BS	CAN	(	8	H	X	h	x
1001	HT	EM	)	9	I	Y	i	y
1010	LF	SUB	*	:	J	Z	j	z
1011	VT	ESC	+	;	K	[	k	{
1100	FF	FS	,	<	L	\	l	|
1101	CR	GS	-	=	M	]	m	}
1110	SO	RS	.	>	N	^	n	~
1111	SI	US	/	?	O	_	o	DEL

Table D.3 The ASCII control characters.

NUL	The null character	DC1	Device control 1
SOH	Start of header	DC2	Device control 2
STX	Start of text	DC3	Device control 3
ETX	End of text	DC4	Device control 4
EOT	End of transmission	NAK	Negative acknowledgment
ENQ	Enquiry	SYN	Synchronization
ACK	Acknowledgment	ETB	End of transmission block
BEL	Bell	CAN	Cancel
BS	Backspace	EM	End of medium
HT	Horizontal tab	SUB	Substitution
LF	Line feed	ESC	Escape
VT	Vertical tab	FS	File separator
FF	Form feed	GS	Group separator
CR	Carriage return	RS	Record separator
SO	Shift out	US	Unit separator
SI	Shift in	DEL	Delete
DLE	Data link escape		

D.3 Eight-Bit EBCDIC

In recognition of the limitations of 5-bit codes, a number of 6-bit codes were developed. The 6-bit Transcode from IBM. supports all uppercase English characters, numbers, symbols, and 16 control characters.

When IBM adopted the 8-bit byte for its large mainframes, it also developed a new character code known as EBCDIC (or Extended Binary Coded Decimal Interchange Code). EBCDIC is *extended* in that it was introduced to replace the older 6-bit codes (sometimes known as binary-coded decimal codes). It supports 256 different character codes (i.e., 28), including all upper and lowercase English characters, numbers, symbols, and control characters. There are many gaps in the EBCDIC table where no specific character translation is defined.

EBCDIC is found primarily in large-scale IBM and IBM-compatible computers as well as support equipment such as terminals and printers. Its character assignments are shown in Table D.4.

Table D.4 EBCDIC Values from 0 to 255

	0000	0001	0010	0011	0100	0101	0110	0111	1000	1001	1010	1011	1100	1101	1110	1111
0000	NUL	DLE	DS		SP	&	-						{	}	\	0
0001	SOH	DC1	SOS						a	j	~		A	J		1
0010	STX	DC2	FS	SYN					b	k	s		B	K	S	2
0011	ETX	DC3							c	l	t		C	L	T	3
0100	PF	RES	BYP	PN					d	m	u		D	M	U	4
0101	HT	NL	LF	RS					e	n	v		E	N	V	5
0110	LC	BS	ETB	UC					f	o	w		F	O	W	6
0111	DEL	IL	ESC	EOT					g	p	x		G	P	X	7
1000		CAN							h	q	y		H	Q	Y	8
1001	RLF	EM							i	r	z		I	R	Z	9
1010	SMM	CC	SM		¢	!	|	:								
1011	VT				.	$	,	#								
1100	FF	IFS		DC4	<	*	%	@								
1101	CR	IGS	ENQ	NAK	(	)	_	'								
1110	SO	IRS	ACK		+	;	>	=								
1111	SI	IUS	BEL	SUB	|	¬	?	"								

Appendix E
Single-Port Operations

Much of the Commkit software has been written for PCs that support two serial communication ports. In certain situations, it may be financially impossible to purchase an additional serial port, potentially rendering many of the experiments described in the book unworkable. Fortunately, much of the software can be rewritten to work with a single serial port, the parallel port, or both. In all cases, it is necessary to have a copy of Turbo C to recompile the changes.

The Commkit software unaffected by the use of a single port includes the interprocess communication software, the point-to-point telephone utility, the modem control software, and all of the stop-and-wait protocol. In addition, the wide area network software can be made to function with only two stations, although routing is still an issue.

Suggestions on modifying the remaining software (the line analyzer, the bus network, the ring network, and the bridge) are discussed in this appendix.

E.1
Basic Concepts

There are two problems associated with the execution of software that requires two serial ports:

Waiting for data on a nonexistent serial port. Not surprisingly, if the serial port doesn't exist, it is impossible for the software to receive data from it.

Waiting for an indication that a transmission has completed. Software can write to a nonexistent serial port, but a transmission completion indication will never be returned. Should the software require a transmission completion indication before sending the next character, a deadlock condition will result.

The objective of any software modification is to ensure that neither of the above conditions is encountered.

E.2 The Parallel Port

The parallel port can be used in place of a serial port. For example, instead of communicating through serial port 1 (`SP1IH`), the software can be written to communicate through the parallel port, `PPIH`.

Using the parallel port requires the use of special cabling and the parallel port software, `parlport.c`. A complete description of the necessary cabling and how the parallel port software functions can be found in Chapter 14.

In all situations, the parallel port must connect to another parallel port.

E.3 The Line Analyzer

The line analyzer should have two serial ports in order to function. If only one additional serial port card is available, it should be used in a PC as a line analyzer to allow the monitoring of network traffic and thereby facilitate error detection and correction.

However, the parallel port can be used in place of one of the serial ports. Bytes are received from and sent to `PPIH`. Sanity checks are still required.

E.4 The Bus Network

The bus network software requires the use of two serial ports: when transmitting, bytes are sent out both ports; when receiving, the byte received is forwarded out the "other" port.

Single-port operation can be achieved by ensuring that only one port is accessed (i.e., transmission occurs out a single port only, and bytes are never forwarded during a reception). For example, by setting `bus.x_state` to `SEND_DONE` for serial port 2, transmissions should never occur through this port. Similarly, if `bus.etxfnd` is set to `TRUE` for serial port 2, collisions can be handled when they occur on serial port 1.

Forcing collisions in a single-port configuration may best be achieved if the PCs operate at a very slow line speed. If the line analyzer is available, it should be used: first to allow the collision to be monitored and second to help slow down the rate of transmission.

The parallel port software can be used in the bus network. It may be necessary to modify `pp_protocol()` so that it returns `XMITDONE` and `RECVDONE` status codes before calling the bus emulator. Additional changes to the transmission software are also required for the transmission of bytes out the parallel port. It may be necessary to use the `send()` primitive in order to supply bytes to the parallel port.

E.5
The Ring Network

The ring network software receives bytes from serial port 2 (`SP2IH`) and forwards them on serial port 1 (`SP1IH`). By sending and receiving through the same serial port, the ring network can be made to function between a pair of PCs. The changes required are minimal, since all transmissions are sent via serial port 1 and input can occur on either port (the source of the receive completion interrupt is not checked).

The parallel port software can also be used as part of the ring network. Although most of the ideas discussed with respect to the bus network are applicable to the ring, *two* sets of software must be produced:

One set must be able to transmit on the serial port and receive on the parallel port.

The other must be able to receive on the serial port and transmit on the parallel port.

In addition, as shown in Figure E.1, an even number of PCs is required.

E.6
The Commkit Bridge

Since the Commkit bridge already supports the serial ports and the parallel port, only the network layer software, using the single-port software, is necessary. Messages can be sent from station to station across the network, and the bridge statistics can still be obtained.

E.7
A Final Word

Before proceeding with the changes suggested in this appendix, check the PCs againsometimes the serial ports are not labeled as such. Also, many PCs are now supplied with "multi-function enhancement peripheral cards" (especially PC/ATs). These cards have a parallel printer port, a game port, and *two* serial ports. In some cases, only one serial port is connected. Check whether the PC is supplied with such a card.

And finally, remember to make copies of any software before modifying it!

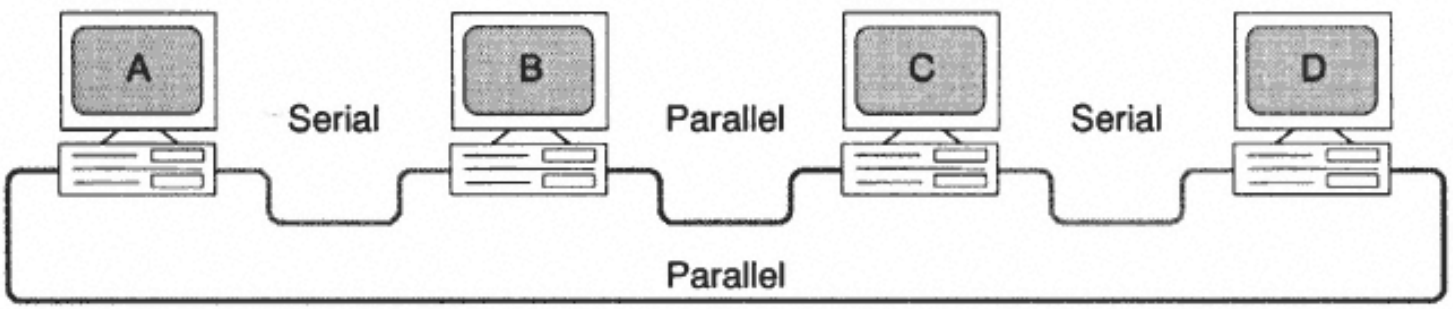

Figure E.1
Creating a ring with serial and parallel ports.

Appendix F
Suggested Readings

Part 1
Background

Lipschutz, Seymour. *Essential Computer Mathematics*. Schaum's Outline Series. New York: McGraw-Hill, 1982.

Sarch, Ray (Ed.). *Basic Guide to Data Communications*. New York: McGraw-Hill, 1985.

Turbo C Reference Guide, Version 2.0. Scotts Valley, Calif.: Borland International, 1988.

Part 2
Point-to-Point Communications

Data Communications, Local Area Networks, UARTs Handbook. Santa Clara: National Semiconductor, 1990.

McNamara, John E., *Technical Aspects of Data Communications* (2nd ed.). Digital Press: Bedford, Mass., 1982.

Technical Reference Manual for the IBM Personal Computer. New York: IBM Corporation, 1983.

Thorne, Michael. *Programming the 8086/8088*. Benjamin/Cummings, 1986.

Part 3
Longer-Distance Communications

General Information. Binary Synchronous Communications. New York: IBM Corporation, GA27-3004-2, 1970.

Datapac: Standard Network Access Protocol Specification. Ottawa: TransCanada Telephone System (now Stentor Communications), 1976.

IBM Synchronous Data Link Control: General Information. New York: IBM Corporation, GA27-3093-2, 1979.

Martin, James. *Security, Accuracy, and Privacy in Computer Systems*. Englewood Cliffs: Prentice-Hall, 1973.

Moshos, George. *Data Communications: Principles and Problems*. St. Paul: West Publishing Company, 1989.

Schwaderer, David. *C Programmers Guide to NETBIOS*. SAMS, 1988.

Sharma, Roshan, Paulo de Sousa, Ashok Ingle. *Network Systems*. New York: Van Nostrand Reinhold, 1982.

Tanenbaum, Andrew. *Computer Networks* (2nd ed.). Englewood Cliffs: Prentice-Hall, 1988.

Part 4
Local Area Networks

Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications. New York: ANSI/IEEE Standard 802.3-1985.

Huchinson, D., J. Mariani, and D. Shepherd (Eds.). *Local Area Networks: An Advanced Course*. New York: Springer-Verlag, 1985.

Keiser, Gerd. *Local Area Networks*. New York: McGraw-Hill, 1989.

Microcommunications Handbook. Santa Clara: Intel Corporation, 1985.

Token-Passing Bus Access Method and Physical Layer Specifications. New York: ANSI/IEEE Standard 802.4-1985.

Token Ring Access Method and Physical Layer Specifications. New York: ANSI/IEEE Standard 802.5-1985.

Part 5
Internetwork Communications

Comer, Douglas. *Internetworking with TCP/IP*. Englewood Cliffs: Prentice Hall, 1988.

Huchinson, D., J. Mariani, and D. Shepherd (Eds.). *Local Area Networks: An Advanced Course*. New York: Springer-Verlag, 1985.

Keiser, Gerd. *Local Area Networks*. New York: McGraw-Hill, 1989.

Sloman, Morris, and Jeff Kramer. *Distributed Systems and Computer Networks*. Englewood Cliffs: Prentice-Hall, 1987.

INDEX

.c

analyzer.c, 123, 489

`clk_ih()`, 125

`clock_handler()`, 126

`do_analyzer()`, 125, 126

`get_char()`, 127

`kb_ih()`, 125

`low_level()`, 125, 127

`sp1_ih()`, 125

`sp2_ih()`, 125

Sanity check, 127

bridge.c, 442, 490

`do_bridge()`, 448

`half_bridge()`, 447

`router()`, 447

busex.c, 342, 349, 352, 353

`do_bus_ex()`, 352

`do_display()`, 352

buslow.c, 342, 349, 417, 446, 448, 489

`bus_emulator()`, 350

`collision_action()`, 351

`low_level()`, 350

`recv_abort()`, 351

recv_byte(), 350
recv_start(), 350
send_byte(), 351
send_collision(), 351
send_start(), 350
commkit.c, 37, 50, 51
clear_scr(), 48
clk_ih(), 39
clock_handler(), 39
display(), 46
do_other_tasks(), 45
dump_queues(), 48
get_line_speed(), 89
initialize(), 33, 39
int2hex(), 47
kb_ih(), 38
keyboard_handler(), 38
keyboard_init(), 38
low_level(), 37
low_level_recv(), 44, 46
low_level_send(), 44, 46
message(), 47
move_cursor(), 47
recv(), 33, 42, 206
register_background(), 34, 50
rtn_to_dos(), 33, 39
send(), 33, 42, 206

`serial_port_stats()`, 92

Interrupt handlers, 35, 39

Processes, 46

UART, 88, 89, 120

UART errors, 92

crcsr.c, 315

`crc_calc()`, 154

crctbl.c

`crc_calc()`, 154

`gen_tbl()`, 154

external.c, 245, 249

`do_modem()`, 248

`do_rmt_scr()`, 249

`keyboard_data()`, 248

`low_level()`, 246, 247

`modem_change()`, 248

ipc.c, 52, 54, 476

`clock_handler()`, 53

`clock_signal()`, 52, 53

`display_screen()`, 52, 54

`do_ipc()`, 52, 53

`keyboard_handler()`, 53

`low_level()`, 53

Background processes, 53

parlport.c, 442, 489

`do_pp_protocol()`, 446

`low_level()`, 446

`pp_init()`, 445

pp_out(), 444

pp_protocol(), 446

pp_recv(), 445

pp_xmit(), 443, 445, 446

start_pp_clock(), 445

State machines, 443

pt2pt.c, 96

do_lcl_scr(), 97

do_pt2pt(), 97

do_rmt_scr(), 97

get_char(), 97

low_level(), 96

system_initialize(), 96

ringex.c, 377, 387

do_display(), 386

do_ring_ex(), 386

ringlow.c, 377, 387, 417, 446, 448, 490

check_for_token(), 385

clear_line(), 385

low_level(), 384

recv_byte(), 386

ring_emulator(), 384

ring_idle(), 385

send_byte(), 385

send_dle(), 385

send_token(), 386

s&wdisk.c, 205, 208

s&wex.c

check_clock(), 193

diagnostic(), 193

do_lcl_msg(), 192
do_rmt_msg(), 192
write_message_number(), 192
s&wlow.c, 205
low_level(), 193, 194
process_message(), 195
recv_protocol(), 194
start_transmission(), 194
xmit_protocol(), 193
sockette.c, 413, 414, 420
bind(), 416, 418
gethostbyname(), 414
getport(), 414
get_net_address(), 414
network_layer(), 417
recvfrom(), 416, 417
release(), 414
sendto(), 415, 416
sockette_init(), 415
wait_event(), 418
sockex.c, 418, 420
do_echo(), 418
do_send_and_recv(), 419
do_time(), 418
process_string(), 419
recvfrom(), 419
sendto(), 419

statmux.c, 269
do_18_clk(), 267
do_l_clk(), 268
do_mux(), 267
low_level(), 268
mux_recv(), 268
mux_send(), 268
wanex.c, 307, 319
do_wan_ex(), 312
update_line(), 313
wanlow.c, 307, 319, 417
data_link_layer(), 315
dl_clock(), 316
dl_r_done(), 316
dl_x_done(), 316
low_level(), 316, 317
physical_layer(), 317
update_ack(), 316
wannet.c, 307, 319
get_vcn_tbl(), 313
handle_call_clearing(), 314
handle_call_request(), 313
handle_data_transfer(), 314
network_layer(), 313
M s&wex.c, 191
M s&wlow.c, 191
.h

ascii.h, 50

busdefs.h, 342

busex.h, 342, 344

    Packet structure, 344

commkit.h

    `queue_entry`, 40

    `struct stack`, 50

devices.h, 50

    Attributes, 46

    UART, 120

    UART Line Control Register, 90

    UART Line Speeds, 89

    UART Modem Control Register, 90

    UART Ports, 89

fcntl.h, 206

general.h, 50

    Process identifiers, 35

Necessary header files, 50

ringdefs.h, 377, 385, 387

s&w.h, 193, 199

sockcode.h, 413, 420

sockette.h, 413, 420

stat.h, 206

wancodes.h, 307

wandef.h, 307

wanstrct.h, 307

A

acknowledgment piggybacking, 170

Active interface, 114

Adaptive routing, 286

ADCCP, 296, 403

Address Resolution Protocol (ARP), 410

Anti-jabber, 339

Application layer, 405

Area code, 212

ARPANET, 409, 440

ASCII, 7, 23, 178, 179

Asynchronous Communication Adaptor, 76

Asynchronous communications, 59, 60, 136

Break condition, 68

Error detection and correction, 136

Errors, 63

Even parity, 64

Frame structure, 63

Framing error, 67

Hardware, 69

Mark parity, 67

No parity, 66

Odd parity, 64

Overrun error, 67

Parity, 69

Parity bit, 64
Parity Calculation, 69
Parity error, 65
Space parity, 67
Start bit, 62
Stop bits, 62
Transmission overheads, 136
Asynchronous Modems, 230
Automatic Calling Units (ACU), 245
Automatic Repeat Request (ARQ), 155
B
Backbone network, 374
Backward Error Recovery, 155
Backward learning, 287
Balanced circuits, 116
Baseband, 333
Baud Rate, 219
Bayonet connector (BNC), 342
Bidirectional Data Transfer, 170
Big endian, 410
Binary Coded Decimal, 8
Biphase, 107
Bisync, 225
Frame Structure, 182
Nontransparent frames, 182
Problems with, 184
Protocol, 178
Synchronization, 184
Transparent frames, 183

Bit, 7
Bit stuffing, 187
Bit time, 74
Bit-oriented protocols, 184
Bits per second (BPS), 219
Bits-per-second (BPS), 62
Block mode, 18
Break condition, 68, 121
Break-Out Box, 122
Bridge Forwarding, 434
Bridge Learning, 434
Bridges, 430
    Source Routing, 436
    Transparent, 431
Broadband, 333
Broadcast Address, 337
Broadcast network, 331
Bus Networks, 331
    1-persistent, 335
    Collision Detection, 334
    Collisions, 334, 358
    CSMA/CD, 335
    IEEE 802.3, 340
    Network Access, 332, 334
    Nonpersistent, 334
    p-persistent, 335
    Token Bus, 358
Bypass relay, 371
Byte, 7, 483

Byte ordering, 8, 410

Byte-oriented protocols, 178

C

Call clearing packet, 289, 290

Call connected packet, 289

Call request packet, 289

Catenet, 393

Central office, 211

Channel, 3

    Duplex, 4

    Simplex, 3

Character codes, 483

    ASCII, 484

    Baudot, 483

    EBCDIC, 486

    IA2, 484

    IA5, 485

Character mode, 18

Cheapernet, 340

Circuit switching, 212

Client-server, 32, 411

Client-servers, 418

Clock, 25

- Initialization, 26
- Latch, 26

Clock drift, 75

Closed-User Group (CUG), 299

Coaxial cable, 109, 280, 336, 340, 361, 430

Codewords, 232

Commkit, 31, 32

- Background processes, 34, 45
- Bridge, 490
- Bus Network, 342, 412, 440, 489
- Clock Interrupts, 38
- CRC Utilities, 153
- External Modem Software, 245
- File Transfer, 205
- Foreground process, 33, 45
- header files, 50
- Internal Modem Software, 250
- Interprocess Communication, 51
- Interrupt Handlers, 35, 45
- Interrupt Vectors, 39
- Keyboard
  - Interrupt Handler, 37
- Layers, 32
- Line Analyzer, 477

Line analyzer, 123, 489
Message handling, 40, 44
Message queue, 40
MS-DOS, 48
Multiplexing, 265
Parallel Port, 489
Point-to-Point Communications, 93
Process, 33
Ring Network, 376, 412, 440, 490
Sockettes, 412
Stop-and-Wait Protocol, 190, 201
Synchronous Software, 189
Transparent Bridge, 440
UART Initialization, 88
UART Interrupt Handler, 91
Wide Area Network, 306, 399, 412
Common Application Service Elements (CASE), 405
Communication
Asynchronous, 15
Full-Duplex, 4
Half-Duplex, 4
Non-transparent, 16
Parallel, 14
Serial, 14
Synchronous, 15
Transparent, 15
Congestion, 284
Connection-oriented services, 291, 438
Connectionless services, 291, 439

Constellation Pattern, 237

Control, 7

Control characters, 16, 183, 485

Control information, 9

Converter

- Parallel-to-serial, 14, 69
- Serial-to-parallel, 14, 62

Core gateways, 440

Correspondent user, 399

CRC-12, 152

CRC-16, 152, 182

CRC-32, 152

CRC-CCITT, 152, 186

Crosstalk, 109

CSMA/CD, 335

Cyclic Redundancy Check (CRC), 147

- Algorithm, 148
- CRC Generators, 151
- Double-Bit Errors, 151
- Error Bursts, 152
- Example, 149
- Odd Number of Errors, 152
- Single-Bit Errors, 151

D

DARPA Protocol Architecture, 409

Data analyzers, 123

Data Circuit-Terminating Equipment (DCE), 215, 274

Data Communications Equipment (DCE), 215

Data link layer, 292, 403, 430

Data packet, 289

Data qualified bit, 300

Data Switching Exchange (DSE), 297

Data Terminal Equipment (DTE), 16

Datagram, 288, 429

DB-25, 228

DB-9, 228

DDCMP, 306

DECNET, 306

Decoding dictionary, 233

Dedicated lines, 274

Demultiplexing, 258

Descrambler, 236

Destination removal, 367

Deterministic access, 358

Dial tone, 212

Dibit, 220

Dielectric, 109

Differential PSK, 217

Digital Network Architecture (DNA), 401

Discovery packet, 437

Dot notation, 440

Double buffering, 72

Downlink, 110

Duplex channel, 4

Duplicate token situation, 370

Dynamic routing, 286

E

EBCDIC, 7, 23, 178, 179

Echo server, 418

Echoplexing, 18

Elastic buffer, 369

Electromagnetic interference, 11

Encoding dictionary, 232

End office, 211

Error Correcting Codes, 175

Ethernet, 336, 430

- LAN coprocessor, 341
- Packet Structure, 337
- tap, 336
- Transceiver, 336, 341

Even parity, 64

Exponential backoff algorithm, 339

F

Facsimile, 17, 238

Faults, 10
    Crosstalk, 109
    Electromagnetic, 11
    Impulse, 11
    Noise, 11
        Intermodulation, 11
        Thermal, 11
        White, 11
    Signal Attenuation, 11, 74, 109
FAX, 238
    Group 3, 238
    Group 4, 238
Fiber Distributed Data Interface (FDDI), 374
File
    Access, 202
    Systems, 201
    Transfer, 200
    Transfer Protocols, 203
Final trunk, 213
Flooding, 285, 431
Flow control, 12, 169
Forward Error Recovery, 175
Forwarding database, 432
Fragmentation, 429
Frame, 10, 63, 139, 182, 183, 185
Frame Check Sequence (FCS), 186
Framing error, 67, 121
Frequency Division Multiplexing (FDM), 259
Frequency-Shift Keying (FSK), 216

Front end, 262

Full-Duplex, 4

Fully connected network, 281

G

Gateways, 394, 428, 437

Gender mender, 112

Generator polynomial, 148

Global address, 283

Go-Back-N, 160, 165

- acknowledgments, 162
- Error Conditions, 163
- Frame Loss, 164
- Identifiers, 161
- Missing Response, 165
- Retransmissions, 172
- Sequence numbers, 167
- Throughput, 165
- X.25, 300
- X.75, 438

Go-back-N

- SDLC, 187, 301
- wanlow.c, 307

GSTN, 211

Guard band, 259

H

Half-Duplex, 4

Half-gateway, 428

Hamming code, 175

Hayes' modem commands, 244

Hierarchical address, 283

High-Level Data Link Control (HDLC), 152, 297, 403

High-usage trunks, 213

Hop count, 285

I

IBM, 486

IBM Token Ring, 371

    Source Routing Bridge, 436

    Synchronous Mode, 373

Icons, 17

Identifiers, 10

IEEE 802, 405

IEEE 802.1, 406, 431

IEEE 802.2, 406

IEEE 802.3, 340, 403, 406, 407, 426

IEEE 802.4, 361, 403, 406

IEEE 802.5, 373, 403, 406, 407, 426

IEEE 802.6, 406

Information frame, 185

Integrated Services Digital Network (ISDN), 111, 212, 403

Intel 8086, 21

Devices, 22

Port numbers, 22

Ports, 22

Intel 8250, 76, 121

Intel 8251A, 188

Intel 8253, 25

Intel 8259, 29, 86

Interrupt mask, 29

Intel 8273, 188

Intelligent bridge, 433

Internet, 393

FTP, 412

HTTP, 412

Rlogin, 411

Routing, 439

SMTP, 411

TELNET, 412

Internetwork Communications, 393

Interprocess Communication (IPC), 31

Interrupt, 28

8086, 28

Handler, 28

Identification Register, 87

Number, 28

Vector, 28, 87

Interrupt Identification Register, 87

Intranetwork Communications, 393

K

Kernel, 31

Keyboard, 23

- Algorithm, 24
- Break scan code, 23
- Make scan code, 23
- Mapping Tables, 25
- Scan code, 23

L

LAP-B, 297

LAP-M, 235

Layered architecture, 399

Layering, 6, 291, 395, 398

Leased lines, 274

Level Converter, 104

Limited-distance modems, 230

Line analyzers, 123

Line Speed, 62

Line Turnaround, 4

Link Access Procedure-Balanced (LAP-B), 297, 403

Link Access Procedure-Modems (LAP-M), 235

Little endian, 410

Local area networks, 330

Local Area Networks (LAN), 330

Local echo, 18

Logical Link Control (LLC), 407

Longitudinal Redundancy Check (LRC), 144

Loopback testing, 120

M

Mailbox, 398

Make

analyzer.exe, 128

busb.exe, 449

busex.exe, 352

bussox.exe, 420

external.exe, 249

ipc.exe, 54, 476

ringb.exe, 449

ringex.exe, 387

ringsox.exe, 420

s&wex.exe, 196

statmux.exe, 269

wanex.exe, 319

wansox.exe, 420

Manchester encoding, 108, 142, 337

Manufacturing Automation Protocol (MAP), 361, 405

Mark, 61, 116, 215, 216

Media Access Control (MAC) Sublayer, 340, 361, 406, 430

Mesh topology, 281

Message, 10, 31, 137, 288

Messaging systems, 288

Metropolitan Area Network (MAN), 374

Minipacket, 375

Missing token situation, 368

Modem, 215

Amplitude Modulation (AM), 215

Answering, 215, 224

Asynchronous, 230
Bell, 212, 221
CB-CF Common, 226
CB-CF Separation, 226
Clear to Send, 225
Frequency Modulation, 216
Full-Duplex Communications, 223
Half-Duplex Communications, 225
Limited-Distance, 230
Originating, 215, 223
Phase Modulation (PM), 216
Request to Send, 225
Synchronous, 230
Modem Status Change Interrupt, 88
Modem Testing, 242
Analog Loopback Test, 242
Digital Loopback Test, 242
Modified Digital Loopback, 243

Remote Loopback Test, 243

Self Tests, 242

Monitor, 31

Monitor passed bit, 369

Monitor station, 368

MS-DOS, 38, 48, 51

chaining, 39

Clock interrupt handler, 39

Interrupt Handlers, 39

Multicast Address, 337, 355

Multipath routing, 286

Multiplexer, 258

Multiplexing, 256

Mux, 258

N

N-protocol, 399

N-services, 399

Name server, 424

Negative acknowledgment, 156

Network analyzers, 356

Network layer, 292, 403, 437

Networks

Bus Networks, 331

Interconnected, 426

Local Area, 329

Metropolitan Area Networks, 358

Packet Switching, 284

Ring Networks, 358

Store-and-forward, 288

Neutral working, 114

Nibble, 7

Non-Return to Zero

Inverted (NRZ-I), 106

Level (NRZ-L), 105

Non-Return to Zero (NRZ), 105

Mark (NRZ-M), 106

Space (NRZ-S), 106

Non-transparent, 182

Non-transparent communication, 16

Nr, 162, 170, 235

Ns, 162, 170, 235

Null-modem, 113, 228

Nyquist, 214

O

Octet, 298

Odd parity, 64

Off-hook, 212

On-hook, 212

Open Systems Interconnection (OSI) Reference Model, 402

Operating system, 31

Optical fiber, 110

Overrun error, 67, 121

P

Packet, 10, 283, 331, 337, 437

Packet Assembler/Disassembler (PAD), 295, 304, 438

Packet Filtering, 432

Packet Switching Exchange (PSE), 284

Packet Switching Networks (PSN), 284

Packet Switching Node (PSN), 284

PAD profile, 305

Parallel port, 442, 489

Parity, 145

Parity bit, 64

Parity error, 65, 121

Passive interface, 114

Peer entities, 293, 399

Permanent virtual circuit, 297

Phase Coherent Phase-Shift Keying (PSK), 217

Physical Connections, 103

Physical layer, 292, 403, 430

Piggybacking, 170

Polar working, 114

Polling, 27

Polling loop, 27

Port Selectors, 262

Positive acknowledgment, 155

Presentation layer, 404

Private lines, 274

Processes, 31

Promiscuous mode, 338, 431

    buslow.c, 446

    ringlow.c, 386, 446

Protocol Concepts, 6

    Channel Utilization, 13

Coding of Information, 6

Control Information, 9

Error Checking and Recover, 10

Synchronization and Timing, 14

Transparency, 15

Protocol Control Information (PCI), 400

Protocol converter, 428

Protocol Data Units (PDU), 399

Protocol entity, 399

Public Data Networks (PDN), 274

Q

Quarantining, 404

R

Reassembly, 429

Receive Line Status Interrupt, 88

Received Data Available Interrupt, 88

Receiver, 4

Register Insertion Ring, 376

Remote echo, 18

Remote Job Entry (RJE), 178

Remote Procedure Calls (RPC), 425

Repeaters, 430

Reset packet, 302

Restart packet, 302

Return to Zero (RZ), 107

Reverse path, 308

Reverse path forwarding, 451

Ring Networks, 358, 363

- Errors
  - Circulating Busy, 368
  - Duplicate Token, 370
  - Missing Token, 368
  - Monitor Failure, 370
- Monitor Station, 368
- Multiple Token, 367
- Register Insertion, 376
- Single Frame, 366
- Single Token, 367
- Token ring, 364

Router, 437

Routing

- Algorithm, 284
- Dynamic, 286
- IEEE 802, 405
- Internet, 439
- Source routing bridges, 431, 436
- Strategy, 284
- Table, 285

X.25, 403
X.75, 438
Routing tables, 439
M Gateways, 429
RS-232-C, 111, 221, 231, 296, 403
25-pin Connectors, 112
9-pin Connectors, 228
Connections, 222
Limitations, 134
Signal Levels, 113
RS-366, 245
RS-422-A, 115, 230
RS-423-A, 115, 230, 276
RS-449, 115, 230, 296, 403
S
Sanity check, 102, 127
Satellites, 110
Scrambler, 236
SDLC, 152, 184, 403
FCS, 186
Frame Structure, 184
Information frame, 185
Poll/Final bit, 186
Supervisory Frame, 185
Transparency, 187
Unnumbered Frame, 185
Segment, 336, 430
Selective Retransmission, 172
Sequence number, 12, 157

Serial interface, 76, 341

Serial port, 76

Service access points (SAP), 400

Service Data Unit (SDU), 400

Service provider, 399

Service user, 399

Session layer, 404

Shift register, 69

Simplex channel, 3

Sliding Window Protocol, 167

    Sequence numbers, 170

Slotted Ring, 375

Socket, 397, 411

Source address, 283

Source removal, 367

Source routing, 436

Space, 61, 116, 215, 216

Spanning tree, 435

Specific Application Service Elements (SASE), 405

Standards Organizations

    American National Standards Institute (ANSI), 485

    Comite Consultatif International Telephonique et Telegraphique (CCITT) CCITT), 111

    Electrical Industry Association (EIA), 111

    Institute of Electrical and Electronic Engineers (IEEE), 111

    International Organization for Standardization (ISO), 111, 402

    International Telecommunications Union (ITU), 111

Star topology, 281

Start bit, 62, 75

Statistical Multiplexing, 261

Stick parity, 79

Stop bits, 62

Stop-and-wait, 155

- Algorithms, 156
- Duplicate Frames, 157
- Receive State Machine, 160
- Retransmission, 156
- Sequence numbers, 159, 166
- Transmission State Machine, 159

Switch, 262

Switch virtual circuit, 297

Synchronous Communications, 136

- Error detection, 142

Error Recovery, 154

Frame Delimiting, 137

Frame Structure, 139

Hardware, 187

Reducing Transmission Overheads, 137

Timing, 140

Synchronous Modems, 230

Synchronous Protocols, 178

System Network Architecture (SNA), 306, 401

T

Technical and Office Protocols (TOP), 361

Terminal concentrator, 262

Terminals

Block mode, 18

Character mode, 18

Terminator, 332

Throughput, 13

Time Division Multiplexing (TDM), 214, 260

Time server, 418

Token, 359

Token ring, 364

Transceiver, 336, 341

Transmission Control Protocol (TCP/IP), 411, 439

Transmission Media, 4, 108

Bounded, 108

Coaxial Cable, 109

Microwave, 110

Optical Fiber, 110
Satellites, 110
Twisted Pair, 109
Two-wire open lines, 108
Unbounded, 108
Transmission window, 167
Transmit Holding Register Empty, 88
Transmitter, 4
Transparent bridges, 431
Transparent communication, 15
Transponder, 110
Transport layer, 403
Trellis Encoder, 237
Tribit, 236
Trunks, 213
Turbo Assembler, 54
Turbo C, 33, 54, 206, 455
`close()`, 206
`getvect()`, 36
`inportb()`, 22
`interrupt`, 36
`make` utility, 475
`open()`, 206
`outport()`, 22
`outportb()`, 22
`read()`, 206
`setvect()`, 36
`tcc`, 476

`tlink`, 476

`write()`, 206

Turbo Linker, 55, 476

20-Milliampere Current Loop, 114

Twisted pair, 109

Two-wire local loop, 212

Two-wire open lines, 108

U

UART, 69, 76, 119

Break condition, 74

Break Detection, 85

Divisor Latch Access Bit, 79, 80

Double buffering, 72

Error Handling, 85

Errors, 84

Even Parity Select, 79

Framing Error, 85

Initialization, 77, 84, 87

Internal overrun, 82

Interrupt Enable Register, 86, 241

Interrupt Identification Register, 87, 241

Interrupt Testing, 121

Interrupt Vectors, 87

Interrupts, 76, 86

Line Control Register, 77, 79

Line Speeds, 78

Line Status Register, 82, 83, 121

Loopback testing, 120

Modem, 239

Modem Control Register, 120, 239

Modem Interrupts, 241

Modem Status Change Interrupt, 88

Modem Status Register, 240

Overrun Error, 73, 84

Parity Enable, 79

Parity Error, 74, 85

Programming, 77

Receive Buffer Register, 83

Receive Line Status Interrupt, 88

Receive Polling, 84

Received Data Available Interrupt, 88

Reception, 72, 83

    Holding Register, 73

    Parity, 72

Sampling clock, 72

Set Break, 80

Start bit detection, 75

Stick Parity, 79

Stop bits, 79

Transmission, 69, 81

Transmit Holding Register, 71

Transmit Holding Register Empty, 82, 88

Transmit Polling, 82

Transmit Shift Register, 69

Transmit Shift Register Empty, 82

Word Length, 79

Unbalanced circuits, 116

Unicast Address, 337

Uplink, 110

USART (Universal Synchronous-Asynchronous Receiver-Transmitter), 188

User Datagram Protocol (UDP), 411, 412, 439

User elements (UE), 405

V.11, 116

V.21, 218

V.24, 111, 221, 231

V.27, 236

V.32*bis*. 237

V.35, 116

V.42, 232, 235

V.42*bis*. 233

Vertical Redundancy Check (VRC), 145

Virtual circuit number, 290

Virtual circuit table, 289

Virtual Circuits, 289, 429

Call clearing, 290

Call establishment, 289

Data transfer, 289

PVC, 297

SVC, 297

Vr, 162, 170

Vs, 162, 170

W

Wide Area Network, 273, 280, 405, 426, 438

Window size, 168

X

X.121, 298, 438

X.2, 276

X.20, 275

X.21, 275, 296, 403

X.21*bis*. 296

X.25, 297, 306, 403, 438

X.26, 276

X.27, 276

X.28, 305, 405

X.29, 305, 405

X.3, 304

X.400, 405

X.409, 405

X.75, 438

X.96, 276